(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
4 July 2002 (04.07.2002)

PCT

(10) International Publication Number
WO 02/052864 A2

(51) International Patent Classification[7]: H04N 7/26

(21) International Application Number: PCT/US01/44264

(22) International Filing Date: 27 November 2001 (27.11.2001)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
09/724,015 28 November 2000 (28.11.2000) US

(71) Applicant: SCIENTIFIC-ATLANTA, INC. [US/US]; 5030 Sugarloaf Parkway, Lawrenceville, GA 30044 (US).

(72) Inventor: ROVIRA, Luis, A.; 2643 Whiteleigh Court, Atlanta, GA 30345 (US).

(74) Agents: GARDNER, Kelly, A. et al.; Scientific-Atlanta, Inc., Intellectual Property Department, 5030 Sugarloaf Parkway, Lawrenceville, GA 30044 (US).

(81) Designated State *(national)*: CA.

(84) Designated States *(regional)*: European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR).

Published:
— *without international search report and to be republished upon receipt of that report*
— *entirely in electronic form (except for this front page) and available upon request from the International Bureau*

(54) Title: MPEG IN A LABEL SWITCHING ENVIRONMENT

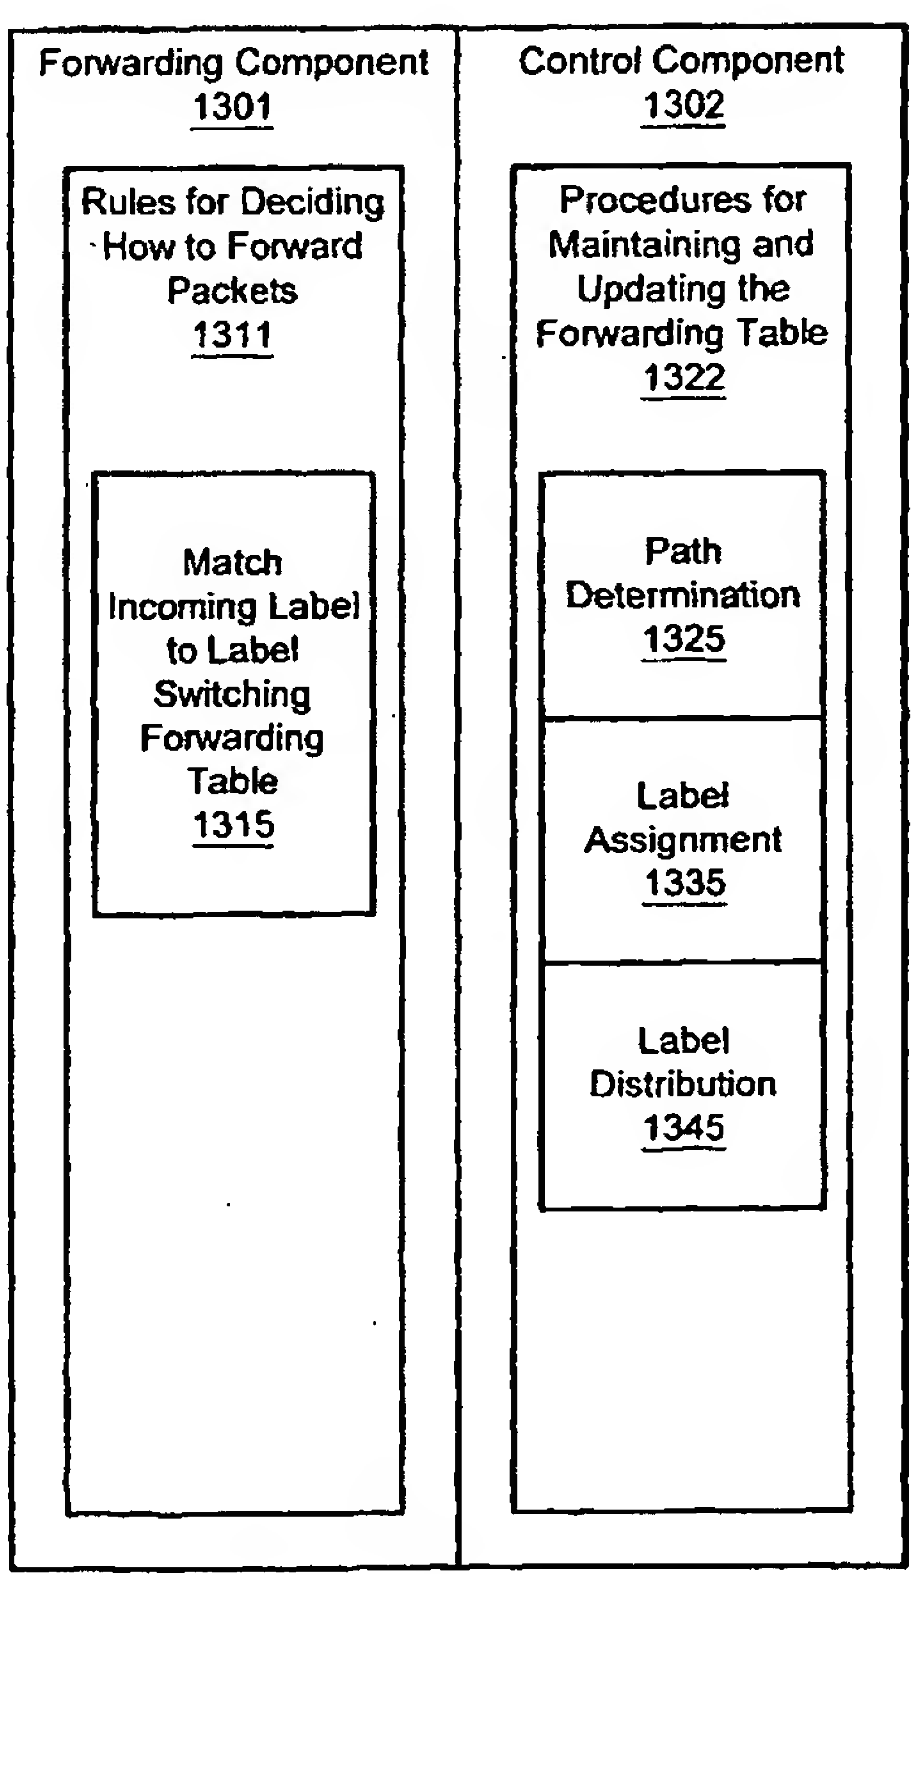

(57) Abstract: The encapsulation of MPEG (Motion Pictures Experts Group) within a label switching protocol such as MPLS (Multiple Protocol Label Switching) is used to implement a multimedia communications network. The network advantageously improves the delivery of customized services to network subscribers. Encapsulating MPEG within MPLS simplifies network deployment by converging the previously distinct environments of data communications and broadcast television service over cable TV networks. In addition, the encapsulation of MPEG within MPLS allows merging MPEG program content streams with IP streams at the MPLS multiplexing level. This allows for applications such as merging MPEG video programs with IP advertising. Furthermore, the MPEG Digital Storage Medium Command and Control (DSM-CC) protocol may be used for label distribution for the label switching network.

WO 02/052864 A2

BEST AVAILABLE COPY

MPEG IN A LABEL SWITCHING ENVIRONMENT

FIELD OF THE INVENTION

The present invention relates generally to the arts of communication networks and multimedia, and more particularly, to a multimedia program delivery system (and associated methodology) for delivering information flows to network subscribers or users.

BACKGROUND OF THE INVENTION

Current digital multimedia networks can generally be divided into a few main categories. One category of digital multimedia networks is a data network based upon the Internet Protocol (IP). IP networks were primarily designed to carry unicast information between two hosts. Though IP has the capability for multicast, IP routers and switches have not been very efficient at forwarding broadcast or multicast information. In fact, the IP network layer is designed to control traffic and prevent many broadcasts from transversing the routers of a network. When IP was developed, the bandwidth of switching devices and routers far exceeded the bandwidth of transport and transmission technologies. Thus, the IP network layer was designed to make routing decisions to keep information from being broadcast across the slow modem and X.25 links of the original Internet. With the development of fiber optics and the introduction of technologies such as Wave Division Multiplexing (WDM), transport bandwidth has grown at a faster rate than the silicon bandwidth of switching technologies such as routers. As a result it may now be more efficient to allow larger percentages of data to pass through the transport network without the normal access restrictions implemented through network protocols such as IP.

Another category of networks that could be used for delivering digital multimedia to subscribers is a network that creates a connection to each destination. This could be done with point-to-point links, circuit-switched connections, or virtual-circuit, packet-switched connections. Often these connection-oriented technologies allow for network resource allocation to implement quality of service (QoS) for real-time information flows such as multimedia audio or video programs. However, these connection-oriented technologies become unreasonable as the number of subscribers needing connections

increases. Thus, these connection-oriented technologies do not scale very well. Furthermore, there have been many difficulties encountered in using connection-oriented ATM transport networks to deliver connectionless IP datagrams.

Another form of digital multimedia delivery network is the CATV (Community Antenna TeleVision) network or cable TV network where content is generally broadcast or narrowcast to a large number of subscribers on a distribution network. Historically, the CATV network has included many frequency-division multiplexed channels that are distributed to a large number of subscribers. In the past, the CATV network has had a large amount of bandwidth to deliver information or data to many subscribers, but the CATV network's capabilities were limited for delivering customized programming for specific subscribers. This made it difficult to deliver subscriber-controllable and network-controllable features through existing CATV networks.

Switching and multiplexing technologies are well known in the art, but a brief description is included in the following paragraphs for clarity. Those skilled in the art will be aware that spatial separation, time separation, and frequency separation are three common ways of separating electromagnetic signals to allow multiple electromagnetic signals to be communicated. The concept of a communication media or medium can be thought of as a signal path that is spatially separated from the interference or noise of other interfering signals enough to allow the communication of information with the signal. This signal path may be a physical link or a wireless link. Within a communications medium time-division multiplexing and/or frequency-division multiplexing are often used to allow the communications medium to carry multiple information flows. Furthermore, the various multiplexing methods can be used within other multiplexing methods. For example, time-division multiplexing can be used inside of frequency-division multiplexing. This is a non-limiting example as there are many ways of combining multiplexing methods and creating multiplexing hierarchies of the same type of multiplexing method.

Digital signals may be communicated on a communications medium using various methods. One simple method is to use a square wave to communicate the signals. Networks using such square waves are often referred to as baseband or unmodulated (*i.e.*, non-frequency-shifted) transmission systems. These networks may allow multiple signals to share a communications medium through time-division multiplexing. This is in contrast to modulated signals that are frequency shifted outside of the base bandwidth (*i.e.*, the baseband) of a square wave representation of the data. Modulation may also be

used to share a communications medium with multiple signals, each signal being placed within its own frequency band. This technique is known as frequency-division multiplexing. In some of the original IEEE (Institute for Electrical and Electronic Engineers) LAN (Local Area Network) standards, the terms baseband and broadband were used to indicate unmodulated, baseband transmission and modulated transmission, respectively. For instance for Ethernet, 10BaseT means 10 MBPS (mega-bits per second), baseband, twisted-pair while 10Broad36 described a frequency-shifting, 10 MBPS, broadband system. The term broadband also has been used to describe communication systems with relatively higher capacity bandwidth as opposed to systems with relatively lower capacity bandwidth that are often known as narrowband systems.

The two primary types of time-division multiplexing are 1) static or fixed TDM and 2) dynamic or statistical TDM. Static or fixed TDM is used in circuit-switching technologies such as the PSTN (Public Switched Telephone Network) as built in the 1960s to 1990s. In static or fixed TDM, once the capacity of a time channel is allocated for a circuit-switched connection, the bandwidth is used up regardless of the actual transmission of information over the time channel. Dynamic or statistical TDM is used in packet-switching to more efficiently allocate bandwidth based on an as-needed or on-demand basis. Common packet-switching technologies include Ethernet switching, Internet Protocol (IP) routing, X.25, frame relay, and ATM (Asynchronous Transfer Mode).

Packet-switching technologies can be categorized into two types: 1) connectionless or datagram packet-switching and 2) connection-oriented or virtual-circuit packet-switching. In connectionless or datagram packet-switching, no connections or paths are established prior to the source being able to communicate with the destination. Instead, the packet switch or router forwards each packet based on a path or route determined at the time that the packet switch or router receives the packet. In datagram packet-switching networks, each packet transmitted from the source to the destination may follow a different path through the network. Due to different delays from following different paths, the packets in a datagram packet-switching network may arrive at the destination in a completely different order than they were transmitted by the source. IP is a common example of a connectionless, datagram packet-switched protocol.

In contrast, in connection-oriented or virtual-circuit packet-switching a connection or virtual circuit is setup between the source and destination that want to communicate. The connection allocates resources in the network between the source and destination,

and the connection establishes the network path to be used by packets communicated from the source to the destination. In virtual-circuit packet-switching the packets transmitted from the source to the destination all generally follow the originally allocated path and arrive at the destination in the same order as they were transmitted by the source. Common examples of packet-switching technologies that use connection-oriented, virtual-circuits are X.25, frame relay, and ATM. There are two main types of virtual circuits: 1) permanent virtual circuits (PVCs), which are usually configured by network administrators because they are infrequently altered and 2) switched virtual circuits (SVCs), which are often setup on demand by network users and applications. SVCs are generally more complicated because they require protocols to allow network users and applications to signal the network for establishing, managing, and releasing the SVCs.

Due to the historical limitations inherited from many of the earlier networking technologies, it has often been difficult, inefficient, and costly to deploy networks capable of broadcasting and multicasting digital multimedia information to subscribers as well as providing data connectivity to these subscribers. These historical networks usually encapsulated digital multimedia packets such as those formatted to the MPEG (Motion Pictures Experts Group) standards into packets conforming to network level protocols such as the Internet Protocol (IP). For IP, these network level packets are known as IP datagrams. Furthermore, advances in label switching technology, as found in such protocols as MPLS (Multi-Protocol Label Switching), have often focused on encapsulating network level packets such as IP datagrams into label switched frames. MPLS was primarily designed to carry network protocols with a special emphasis on the Internet Protocol (IP) due to the ubiquitous nature of IP based networks. Also, the common method of delivering MPEG packets is to encapsulate them inside IP datagrams. However, a simple straight-forward combination of label switching technologies to encapsulate network level packets such as IP datagrams that then encapsulate MPEG packets leads to inefficiencies and makes it more difficult to deliver broadcast and multicast services to subscribers. This problem in providing digital services that are similar to the services provided by CATV broadcast networks is at least partially due to the addressing limitations of the IP protocol.

Though IP is a robust protocol that offers many advantages, it was designed at a time when network technologies were significantly different than they are today. Due to the large deployment of IP capable equipment and the large number of networks carrying IP traffic, many network improvements focus on the IP protocol without ever

reconsidering some of the initial design decisions made in the IP protocol. When IP was designed, the bandwidth of wide-area links between computers was severely limited. As a result IP routers generally limit traffic on wide-area links by filtering packets based on IP addresses. With the increased bandwidth of fiber optic cable, these assumptions of the Internet protocol (IP) may no longer be valid. Furthermore, the filtering rules of IP routers may actually be limiting the simplified deployment of some services such as broadcast and multicast digital multimedia. Thus, this disclosure suggests solutions to these networking problems and reconsiders the necessity for carrying MPEG packets within IP datagrams inside of MPLS frames.

SUMMARY OF THE INVENTION

The preferred embodiments of the present invention provide a method (and associated devices) of using MPEG packets encapsulated in a label switched protocol such as MPLS (Multiple Protocol Label Switching) to overcome limitations of previous networks. An embodiment of the present invention includes implementing label switching functionality, implementing MPEG functionality, and encapsulating MPEG packets within label switching functionality. Normally, MPEG would not be considered a network protocol, which is what label switching protocols are generally designed to encapsulate. However, the use of the preferred embodiments of the present invention in a network advantageously improves the delivery of customized services to network subscribers. Encapsulating MPEG within MPLS simplifies network deployment by converging the previously distinct environments of data communications and broadcast television service over cable TV networks. In addition, the encapsulation of MPEG within MPLS allows merging MPEG program content streams with IP streams at the MPLS multiplexing level. This allows for applications such as merging MPEG video programs with IP advertising, as one non-limiting example. Furthermore, for those label switching devices that are MPEG-aware, the MPEG Digital Storage Medium Command and Control (DSM-CC) protocol may be used for label distribution for the label switching network. Other common protocols may be used to distribute labels to non-MPEG-aware label switching equipment. Thus, all the label switching equipment in the network does not have to be MPEG-aware.

The preferred embodiments of the present invention advantageously simplify the delivery of multimedia information flows to network subscribers and promote the

convergence of previously distinct historical perspectives on network technologies. These historical perspectives primarily developed from the data networks for computers that predominately use unicast and multicast packets and from the CATV networks for video and other multimedia content delivery that generally use broadcast and narrowcast. The preferred embodiments of the present invention resolve some of the previous problems with the delivery of digital multimedia services to subscribers as well as allow simplified deployment of new services such as personal television.

The preferred embodiments of the present invention advantageously use MPEG packets in a label switching environment to allow the convergence of many of the previously described network technologies. The preferred embodiments of the present invention overcome many of the limitations of the previously described network technologies to deliver multimedia content to subscribers. In an embodiment of the present invention, the network is designed based on the understanding that the restrictions of an IP network layer are not well suited to the deployment of a multimedia network where much of the traffic is broadcast or multicast to many subscribers. Furthermore, the increased transport bandwidth of fiber optic cables has made an IP network layer less of a necessity. Also, the preferred embodiments of the present invention allow for service flow or information flow resource allocation in the network to provide for quality of service (QoS) in multimedia delivery without all the complexity of end-to-end, connection-oriented technologies such as ATM. In addition, the preferred embodiments of the present invention more easily allow the multiplexing of subscriber-customized information within the network than was previously possible in other existing CATV networks. This subscriber-customized information could be used to implement new services.

As stated above, although the preferred embodiments of the present invention are oriented to time-division multiplexing techniques, they may be used within one or more frequency channels on a frequency-division multiplexed medium such as current CATV networks. Alternatively, the preferred embodiments of the present invention could be used on networks without frequency-division multiplexing. Though the current cable TV network has a broadband structure that uses frequency-division multiplexing for broadcast channels, the preferred embodiments of the present invention could be used to carry these broadcast channels in addition to subscriber-customized information in a time-division multiplexed format. Thus, the preferred embodiments of the present invention may allow a migration away from the broadband, frequency-division multiplexed

architecture of current cable TV networks towards a newer baseband, time-division multiplexed architecture for cable TV networks. This migration may occur in stages, as the preferred embodiments of the present invention will likely first be used to provide additional and new services in a time-division-multiplexed format. These example uses of the preferred embodiments of the present invention are not intended to be limiting, and those skilled in the art will be aware of unlimited variations for combining the preferred embodiments of the present invention with various time-division and frequency-division multiplexing techniques. This disclosure is generally not meant to imply any limitations as to combining the preferred embodiments of the present invention, which is generally related to time-division multiplexing and baseband communications, with other multiplexing or transmission techniques.

BRIEF DESCRIPTION OF THE DRAWINGS

The invention can be better understood with reference to the following drawings. The components in the drawings are not necessarily to scale, emphasis instead being placed upon clearly illustrating the principles of the present invention. Moreover, in the drawings, like reference numerals designate corresponding parts throughout the several views. The reference numbers in the drawings have at least three digits with the two rightmost digits being reference numbers within a figure. The digits to the left of those two digits are the number of the figure in which the item identified by the reference number first appears. For example, an item with reference number 211 first appears in FIG. 2.

FIG. 1 is a diagram of the Open Systems Interconnection (OSI) protocol model;

FIG. 2 is a diagram showing a first location where label switching information may be placed in the OSI model;

FIG. 3 is a diagram showing a second location where label switching information may be placed in the OSI model;

FIG. 4 is a diagram showing label switching information between some example data link protocols and network protocols;

FIG. 5 is a first diagram showing a protocol layer example with the PPP protocol used in a virtual private networking tunnel;

FIG. 6 is a second diagram showing a protocol layer example with the PPP protocol used in a virtual private networking tunnel;

FIG. 7 is a diagram showing MPEG transport inside an IP datagram in a label switched frame where the label is included in a shim label header;

FIG. 8 is a diagram showing MPEG transport inside an IP datagram in a label switched frame where the label is included in the layer 2 address fields;

FIG. 9 is a diagram showing MPEG transport inside an ATM cell or a frame relay packet as used on the PVCs and SVCs of ATM and frame relay;

FIG. 10 is a diagram showing MPEG transport directly inside a label switched frame where the label is included in a shim label header;

FIG. 11 is a diagram showing MPEG transport directly inside a label switched frame where the label is included in the layer 2 address fields;

FIG. 12 is a diagram showing the functions of a standard, layer 3 router;

FIG. 13 is a diagram showing the functions of a label switching router (LSR);

FIG. 14 is a diagram showing the functions of the MPEG Systems Specification;

FIG. 15 is a diagram showing how MPEG elementary streams are encapsulated in Packetized Elementary Streams (PES) and further encapsulated in MPEG transport packets;

FIG. 16 is a diagram showing an example of how MPEG transport uses Packet IDs (PIDs) to multiplex many programs including many Packetized Elementary Streams (PES); and

FIG. 17 is a diagram showing the use of Label Switching Routers (LSRs) in a label switching network for delivering MPEG to subscriber devices.

DETAILED DESCRIPTION

Protocol Layering

Protocol layering allows designers to develop software and hardware without worrying about the details of the implementation of other parts of communication systems. For example, the TCP/IP protocol suite used in the Internet comprises many layers. The TCP/IP protocol suite is named after the Transport Control Protocol (TCP) and the Internet Protocol (IP), which are two of the many protocols in the protocol suite. FIG. 1 is a block diagram illustrating the seven-layer OSI (Open Systems

Interconnection) reference architecture or protocol model for communication systems. The seven layers of the OSI reference architecture are: 1) the physical layer 101, 2) the data link layer 102, 3) the network layer 103, 4) the transport layer 104, 5) the session layer 105, 6) the presentation layer 106, and 7) the application layer 107. The OSI model was developed for an OSI protocol that was not widely accepted by the communications industry. However, the seven-layer OSI protocol model is a useful abstraction for evaluating and discussing communication protocols and has become well known in the art for such purposes. Thus, a brief description of the model will be helpful before describing various embodiments of the present invention. Because the OSI model is well known in the art, a detailed discussion of all the features and functionality of each level in the OSI model will not be covered. Thus, the following descriptions of each layer are necessarily just a high-level overview of the some the functions generally found in various levels of protocol stacks or suites.

As shown in FIG. 1, the first layer of the OSI protocol model is the physical layer 101, which is usually concerned with the communication of raw bits over a transmission channel. This usually involves the electrical or optical signal levels used to represent zeros and ones on the communications medium. Also, the physical layer 101 includes repeaters that regenerate and propagate digital signals. The next level of the OSI model is the data link layer 102, which generally operates to ensure that bits are delivered on the medium without error. The IEEE (Institute for Electrical and Electronic Engineers) divided the data link layer 102 into two further sublayers, the media access control (MAC) sublayer 102a and the logical link control (LLC) sublayer 102b.

The MAC sublayer 102a comprises functions for grouping the bits of information into frames. The frames usually contain headers and trailers (or tails) as well as some form of error checking such as error detection or error correction codes. In addition, MAC frame headers often contain layer 2 addresses or MAC addresses. Layer 2 addresses are usually included in MAC frame headers when multiple streams of information are multiplexed onto a communication channel and the identification of the various streams for demultiplexing will be based on information in the stream. For instance, in many LAN technologies multiple transmitting and multiple receiving devices can use MAC addresses to uniquely identify which devices transmitted frames and which devices are to receive frames on a communications medium that is shared by the multiple devices.

In fact, the Media Access Control 102a sublayer often contains rules for specifying which device of multiple devices can transmit or receive at a specific time and/or frequency on the shared medium. A communications medium shared by a number of devices contending for access is often called a contention medium, a broadcast medium, or a multi-point medium. It is often called a broadcast medium because even though only one device can transmit on the medium at a specific time and frequency, many devices could potentially receive a broadcast frame on the medium at a specific time and frequency. Algorithms for arbitrating or controlling access to a shared medium can be distributed among the devices sharing the medium as is the case in Ethernet or the IEEE 802.3 CSMA/CD (Carrier Sense-Multiple Access with Collision Detection) algorithm. Alternatively a controller can execute a centralized algorithm to allocate transmission rights to resolve contention in broadcast or shared media.

In contrast to multi-point, contention, shared, or broadcast media, there are no problems with conflicts over the use of the media in point-to-point media. A point-to-point medium is a medium connected between only two devices. In a uni-directional point-to-point link, a transmitter in a first device communicates in a forward direction with a receiver in a second device. In a bi-directional point-to-point link, a transmitter in a first device communicates in a forward direction with a receiver in a second device, and a transmitter in the second device communicates in a reverse direction with a receiver in the first device.

The LLC sublayer 102b of FIG. 1 provides support for a selection between a datagram service, a connection-oriented service, and an acknowledged datagram service. The common IEEE protocol for LLC sublayer 102b is 802.2.

Referring now to network layer 103 of FIG. 1, the network layer 103 generally provides transparent transfer of data from the transport layer 104. The network layer 103 handles many of the details of the underlying data transmission, switching, and routing functions. To handle these functions, the network layer 103 often contains addressing information that determines how data is switched or routed through the network. As mentioned above, the abstractions of the OSI model do not exactly match every working protocol including the TCP/IP protocol suite. Still, the IP protocol is generally considered to be a level 3, network layer protocol 103.

A primary function of the transport layer 104 is to accept data from the session layer 105 and break the data into smaller pieces before passing it to the network layer 103. In addition, the transport layer 104 often verifies that the data arrives properly at the

destination. From the TCP/IP protocol suite, the Transport Control Protocol (TCP) and the User Datagram Protocol (UDP) are often thought of as layer 4 or transport layer protocols 104. TCP is a connection-oriented protocol that ensures in order delivery of the data that the TCP layer presents to higher level protocols. UDP is a connectionless, datagram protocol that does not guarantee delivery of information, but UDP is useful for many simple communications that do not need the overhead of TCP.

The next layer of the OSI model is the session layer 105, which among other things provides two-way simultaneous, two-way alternating, or one-way transfer between two devices. The presentation layer 106 handles data translation, formatting, and syntax selection. Finally, the application layer 107 contains many commonly used functions for distributed applications. Examples of application layer protocols from the TCP/IP protocol suite include Telnet, the File Transfer Protocol (FTP), and the Hypertext Transfer Protocol (HTTP) that is used for retrieving web pages.

Label Switching

Many factors have influenced the development of label switching technology and standards. First, a need existed for performance improvements in IP routing to obtain the price-performance characteristics of switches is one of the factors. In addition, problems encountered in integrating IP with virtual-circuit technologies such as ATM have driven the demand of simplified solutions through label switching. Finally, the need to make more sophisticated routing decisions than the destination-based routing available in IP also has contributed to the demand for label switching. Originally, many vendors developed proprietary label switching technologies under various names including IP switching and Tag switching. In order to standardize some of these protocols, the Internet Engineering Task Force (IETF) established a working group that has developed Internet RFCs (Request for Comments) and draft RFCs with regard to Multi-Protocol Label Switching (MPLS). The preferred embodiments of the present invention are designed to work with any label switching technology that is capable of carrying many protocols, and this includes the MPLS standard as defined by the IETF.

MPLS is a label switching protocol that was designed to work over multiple protocols and to carry multiple protocols within MPLS label switched frames. Generally, MPLS was designed to work over an OSI layer 2, data link protocol and to carry an OSI layer 3, network layer protocol. Though MPLS is designed for multiple protocols, the

main proposed use for MPLS by those skilled in the art is to carry IP datagrams (*i.e.*, an IP network layer). In the embodiments of the present invention, MPLS or other label switching technology may be used to carry MPEG transport packets. MPEG transport is not normally considered to be a network layer protocol by those skilled in the art.

In general, label switching works by assigning relatively short, fixed-length labels to packets. The network equipment used to implement label switching is called a Label Switching Router (LSR). All packets that are to be forwarded in an equivalent method by an LSR are said to be part of a Forwarding Equivalence Class (FEC). One example of an FEC includes a set of packets destined for the same destination address prefix. Another example may be a flow of packets with a common quality of service (QoS) between two end-points. Packets forwarded in an equivalent manner by an LSR are assigned the same label and placed in the same output queue of an LSR for transmission through one of the LSR's output ports. LSRs generally perform the following simplified forwarding algorithm for label switched packets or frames: read the label on incoming packets, determine the FEC by looking up the incoming label in a forwarding table, swap or replace the incoming label in the packet with an outgoing label as determined by the FEC, and place the packet with the outgoing label into an output queue and/or interface for transmission as determined by the FEC. In some cases instead of swapping a label, an LSR may add or push a new label onto a label stack, or an LSR may delete or pop a label from a label stack. The use of a label stack allows a label switched frame or packet to carry one or more labels. This capability of using label stacks allows more flexibility for label switching than the use of just a single label. For example, without limitation, a label stack may be used to carry two labels when a packet is forwarded through a transit routing domain. The use of label stacks as opposed to single labels is well known to those skilled in the art and has been specified in protocol descriptions of tag switching and MPLS.

The path that a labeled frame or packet follows over several hops of LSRs is known as a Label Switched Path (LSP). An LSP is somewhat like a virtual-circuit. However, an LSP is not necessarily between the source and destination end points for a packet. Instead, an LSP starts at an ingress LSR, where a label is first assigned to a packet, and ends at an egress LSR, where a label is removed from a packet.

There are two primary ways that label switched paths (LSPs) can be established in networks. First, each label switching router (LSR) may make independent decisions to assign a label to a forwarding equivalence class (FEC). After making such a decision,

then this LSR propagates the label assignment information to its neighbor LSRs. Neighbor LSRs are the LSRs that are adjacent to an LSR along a label switched path (LSP). This type of label assignment process is known as independent control of LSP establishment because the LSRs are generally making independent decisions on the creation of an LSP. In effect the algorithm to create LSPs is distributed through many LSRs. The other primary method of establishing LSPs is through ordered control. In this method the LSP is usually setup from ingress LSR to egress LSR in an ordered fashion. In ordered control of LSP establishment, a decision to create an LSP generally is not made by each LSR. Instead of the distributed decision to create an LSP in independent control, the decision to create an LSP in ordered control is usually made by a more centralized algorithm. This might be done at ingress and/or egress LSRs. Alternatively, the decision in ordered control might involve an algorithm executing on at least one centralized controller.

Some label switching routers (LSRs) have the capability of selectively forwarding packets received with label switching headers using the forwarding algorithms of label switching technologies, while selectively forwarding packets received without label switching headers using forwarding algorithms not associated with label switching technologies. An LSR with such functionality may revert to slower IP routing decisions when it receives an IP datagram without a label header.

Some LSRs allocate labels dynamically to flows such that an initial number of IP datagrams below a threshold count from a source to a destination may be forwarded by an LSR using standard IP routing functions. Then a label will be dynamically allocated to the flow for all IP datagrams from the source to the destination over the threshold count within a certain amount of time. Further IP datagram traffic between the source and destination will be labeled and will be processed using the faster performance forwarding algorithm of label switching as opposed to the slower performance algorithm of IP routing.

The performance of an LSR using such dynamic label allocation functions is somewhat similar to the performance of caches in memory architectures that store data in faster access locations based on the spatial or temporal locality of the data from a previous memory access. Spatial locality (or spatial closeness) usually implies that unaccessed data is likely to be accessed in the future if it is located near data that has just been accessed. Temporal locality (or time closeness) usually implies that data that has just been accessed is likely to be accessed again. For label switching environments, once

a threshold number of packets or IP datagrams has been communicated between a source and destination, it is likely that more datagrams will be communicated between the source and destination in the future. This is the case in file transfers or web page downloads. Once, this threshold is met, the LSR will expend the processing overhead to create a label for the flow of information between the source and destination so that future transfers of data will utilize the faster forwarding algorithm of label switching instead of the slower forwarding algorithms of routing protocols such as IP. This improved performance of label switching is similar to the improved performance from cache memory architectures because data that has been cached will be handled more quickly than uncached data much as information flows with assigned labels will be handled more quickly than information flows without assigned labels.

When the dynamic binding of labels is modified by the receipt of data to be forwarded in the network such as the receipt of a number of IP datagrams over a threshold count, the binding is known as a data-driven label binding. When the receipt of control information modifies the dynamic binding of labels, the binding is known as control-driven label binding. An LSR does not have to use data-driven label binding or control-driven label binding to the mutual exclusion of the other type of binding. In addition, there are many variations of both data-driven label binding and control-driven label binding.

There are two primary ways to encode the information needed for label switching in a packet. First, the label switching information may be carried in a label header "shim" that is placed between layers 2 and 3 of the OSI model. The shim contains the label switching information and is encapsulated between a layer 2 header and a layer 3 header. Alternatively, the label switching information may be included in the layer 2 header of some protocols such as ATM and frame relay. FIG. 2 shows how label switching information 211 can be placed between the data link layer 102 and the network layer 103. The label switching information 211 usually includes short, fixed length labels to allow simplified and fast forwarding of packets. In addition, for some protocols, such as ATM and frame relay, it is possible to incorporate the label switching information into the layer 2, data link layer addresses as show in FIG. 3, which displays the data link layer including label switching information as item 302. For example, in ATM the label switching information could be included in the Virtual Channel Identifier (VCI) field and/or in the Virtual Path Identifier (VPI) field in the ATM header. In frame relay the label switching information could be included in the Data Link Connection Identifier (DLCI) field.

Though some protocols can incorporate label switching information within the layer 2, data link addresses, it is not required that these protocols incorporate the label switching information within the layer 2 header. Instead, these protocols can also use the shim label header. Thus, the shim label header may be used with protocols such as ATM or frame relay. In addition, in ATM or frame relay some of the label switching information could be included in the layer 2 header and some could be included in a shim label header. Still, for efficiency reasons it is likely that protocols such as ATM and frame relay will use the layer 2 headers to incorporate label switching information. Furthermore, ATM and frame relay are only examples of two protocols that currently exist that are capable of incorporating label information within the layer 2 header of the protocol. Other layer 2 protocols may exist that allow label switching information to be incorporated in the layer 2 headers.

FIG. 4 shows how multiple layer 2, data link protocols such as Ethernet 402a, token ring 402b, ATM 402c, frame relay 402d, and PPP (Point-to-Point Protocol) 402e as well as multiple layer 3, network protocols such as IPv6 (Internet Protocol version 6) 403a, IPv4 (Internet Protocol version 4) 403b, IPX (Internetwork Packet eXchange protocol) 403c, and AppleTalk 403d can be used with a label switching layer 411. IPv4 403b is the common IP protocol used in the Internet, while IPv6 403a is a newer revision of the IP protocol. IPX 403c is a protocol often used on Novell networks, while AppleTalk 403d was developed for Macintosh networks. The data link layer protocols and the network layer protocols of FIG. 4 are only examples. Those skilled in the art will be aware that many more data link and network protocols exist and may be used with label switching technologies such as MPLS. Multi-Protocol Label Switching (MPLS) was designed to be used with both multiple network layer protocols as well as multiple data link layer protocols. Thus, MPLS has been described as being multi-protocol both above and below.

Label switching and MPLS in particular do not fit nicely in the 7-layer OSI model. Label switching is not a layer 2, data link level protocol because it can operate over many layer 2 protocols. In addition, it is not a layer 3, network level protocol because it does not have the common addressing and routing functions of a layer 3 protocol. Even though label switching and MPLS violates the abstractions of the 7-layer OSI model, it still provides useful functionality to networks.

Virtual Private Networks

One must be careful to avoid confusion between a point-to-point medium (discussed earlier) and the Point-to-Point Protocol (PPP) 402e shown as a layer 2, data link protocol in FIG. 4. Although PPP 402e was designed to provide multiple-protocol encapsulation over point-to-point media, the robust features of PPP 402e have been used and expanded to allow its use in many additional environments. Specifically, PPP 402e provides the capability to encapsulate many protocols and has control mechanisms to allow for the negotiation of settings for various other protocols. As a result of these features PPP 402e has been used in many tunneling protocols for virtual-private networks (VPN) and other uses.

FIG. 5 shows a possible protocol stack where PPP 402e is used in a VPN. Often VPNs are designed to run over a public network layer 503a such as the IP network of the Internet. Normally, VPNs carry a private network layer 503b inside of a protocol such as PPP 402e. The private network layer 503b could be any network layer protocol. Normally the PPP 402e frames are carried in an encapsulation protocol layer 511. The Point-to-Point Tunneling Protocol (PPTP) is an example of an encapsulation protocol layer 511. PPTP is based on the Generic Routing Encapsulation (GRE) Protocol.

FIG. 6 shows another possible protocol stack where PPP 402e is used in a VPN. The protocol stack in FIG. 6 includes a public network protocol 603a, a private network protocol 603b, a public transport protocol 604a, and a private transport protocol 604b. A public network protocol 603a such as IP in the Internet and a public transport protocol 604a such as UDP may be used to carry an encapsulation protocol layer 611. Example encapsulation protocol layers 611 that would use the protocol stack of FIG. 6 include the Layer 2 Tunneling Protocol (L2TP) and the Layer 2 Forwarding (L2F) protocol.

Those skilled in the art will be aware that the examples of VPNs shown in FIG. 5 and FIG. 6 are only a small portion of the possible ways of combining protocols stacks for VPNs and other purposes. They are presented in this present application to show that there are at least two places where MPLS or other label switching information can be placed in the protocol stacks. Specifically, in FIG. 5 MPLS or other label switching information can be included between data link layer 102 and public network layer 503a to implement label switching in a public network. In addition, in FIG. 5 label switching information could be used between PPP 402e and private network layer 503b to

implement label switching in a private network that has been tunneled through the public network in a VPN.

Also, in FIG. 6 MPLS or other label switching information can be included between data link layer 102 and public network layer 603a to implement label switching in a public network. In addition, in FIG. 6 label switching information could be used between PPP 402e and private network layer 603b to implement label switching in a private network that has been tunneled through the public network in a VPN. Thus, it is possible for private label switching information (that might exist between PPP 402e and private network layer 503b or 603b) to be tunneled through the Internet. The examples in FIG. 5 and FIG. 6 are intended to suggest that the label switching information normally carried between data link layer 2 and network layer 3 (or in layer 2 header information) may occur not only above or in a data link layer 102 directly above a physical layer 101, but also above or in a layer such as PPP 402e that has many of the functions of a data link layer 102 but may be used in a way that is not directly above a physical layer 101. In any case, the prior use of MPLS by those skilled in the art would likely be limited to carrying IP datagrams or other network level traffic. This would not include MPEG because those skilled in the art generally would not consider MPEG transport to be a layer 3, network level protocol.

MPEG Encapsulation

The concept of encapsulation in building frames or packets for transmission is well known by those skilled in the art, and is important in understanding how frames or packets are generated for transmission and decoded in reception using a layered protocol model such as the OSI 7-layer model. In general, for transmission each higher level protocol passes information down to a lower level protocol. The information is encapsulated by the lower level and further passed down to an even lower level protocol for further encapsulation. This process is repeated until the frame or packet is completely created for transmission. An inverse process of passing information up to higher level protocols is performed in decoding received frames or packets. Depending on the protocol, most protocol layers build frames or packets using a header, a payload, and an optional trailer or tail. Some protocols use fixed length frames or packets while other protocols use variable length frames or packets. For example, 10 MBPS Ethernet is a variable length data link level protocol with frames varying in size from 64 octets to 1518

octets. In contrast, ATM has a 53 octet fixed length cell. In addition, there are various ways of handling protocols capable of carrying variable length payloads with the two most common being the inclusion of a length field in the header for that protocol and the use of some special flag pattern to indicate the end of a variable length frame or packet.

Properly mapping information from higher level protocols into lower level protocols sometimes creates problems when the amount of information from a higher level protocol does not fit the size limitations of the lower level protocol. Various methods are commonly known in the art for breaking large blocks of information into smaller blocks of information for transmission by multiple lower level protocol transmissions. In addition, various methods are commonly known in the art for adding stuffing or filling to small blocks of information to fill up the required minimum sizes of frames or packets. Furthermore, as packets transverse a network with different maximum transfer unit (MTU) sizes, the packets created for a network with a large MTU may have to be fragmented into smaller packets for transmission over networks with a smaller MTU. All these variations as well as others variations that are well known in the art for dealing with encapsulation are intended to be included in the embodiments of the present invention. Thus, figures or explanations showing a lower protocol layer header followed by a payload from a higher protocol layer without any trailers or tails are not meant to indicate that the embodiments of the present invention will not work with protocols that have trailers or tails. In addition, figures or explanations showing methods of mapping large blocks of information into a lower level protocol without regard to the minimum and/or maximum size limitations of the lower level protocol are not intended to limit the embodiments of the present invention to only work when their is an exact fit in the size of the upper level information into the size limits of the lower level information. Therefore, FIGs. 7 - 11 are only intended to be example packets or frames. These figures are not intended to suggest any limitations with respect to methods of encapsulating higher level protocol information into lower level protocol information. For example, the fact that FIGs. 7 - 11 do not show tails or trailers to the frames or packets is not an indication that the embodiments of the present invention cannot work in environments where the frames or packets have tails or trailers.

Furthermore, with respect to FIGs. 7 - 11, any references to a label or a label header do not mean to limit a label header to only include a label. Though the predominate information in a label header is often a label, other information is generally also contained in a label. For example, an MPLS shim header contains a 3 bit

experimental field, a 1 bit stack field, an 8 bit Time-to-Live (TTL) field, and one or more 20 bit labels in a label stack. Also, multiple label shim headers might be used within a single label switched frame for additional flexibility.

For datagram packet data networks, MPEG has generally been encapsulated inside a network layer protocol such as IP. FIG. 7 shows an example of how MPEG may be encapsulated in IP in a label switched frame. The label switched frame would include a header with layer 2 addressing 702, an IP network layer 703, and an MPEG transport packet in the payload of an IP datagram 704. The label 711 for label switching would be carried between the header with layer 2 addressing 702 and the header for the IP network layer 703. Often additional protocols such as the user datagram protocol (UDP) may be used between IP network layer 703 and the MPEG transport packet. Thus, the MPEG transport packet may be encapsulated within other protocols, such as UDP, which are then encapsulated inside the payload of an IP datagram 704.

Similarly for protocols such as ATM or frame relay that have the capability of including labels within the layer 2 header, FIG. 8 shows how MPEG is commonly encapsulated in IP in a label switched frame. The label switched frame would include an ATM or frame relay header incorporating a label within the layer 2 addressing fields 802, an IP network layer 803, and an MPEG transport packet in the payload of an IP datagram 804.

In addition to encapsulating MPEG in an IP datagram, MPEG has been encapsulated in ATM cells for virtual-circuit packet networks. FIG. 9 shows how an ATM or frame relay header 902 might encapsulate an MPEG transport packet in the payload of ATM cells of frame relay frames. For ATM this configuration for MPEG has been used for PVCs or SVCs using a Q.2931 as a layer 3 protocol to establish SVCs. Thus, the layer 2 ATM header encapsulates an MPEG transport packet while Q.2931 is used as a layer 3 protocol with layer 3 addresses to set up the ATM virtual circuits to carry the ATM cells.

Though some protocols such as ATM and frame relay are inherently label swapping or switching protocols, the prior use of MPEG transport over ATM frames was not generalized to work over many label switching protocols. In addition, ATM normally uses fixed PVCs or Q.2931 for SVCs to establish end-to-end virtual circuits between MPEG source and destination devices. Furthermore, the MPEG source and destination devices would usually be connected to ATM switches through point-to-point media.

In contrast, the preferred embodiments of the present invention work over other label switching protocols besides ATM. In general, the preferred embodiments of the present invention work over MPLS (Multi-Protocol Label Switching) or other label switching protocols capable of natively or directly carrying MPEG transport, for example. The multi-protocol below capability of MPLS means that in the preferred embodiments of the present invention MPEG may be carried in MPLS over other layer 2 protocols besides ATM and frame relay. In addition, in the preferred embodiments of the present invention the label switched paths (LSPs) are not necessarily end-to-end connections between MPEG devices as in the case of the PVCs and SVCs of ATM and frame relay. Furthermore, in the preferred embodiments of the present invention the labels in general label switching may be distributed by various protocols described below, while the header addresses (*i.e.*, VCI and VPI) of ATM are fixed for PVCs and setup using Q.2931 for SVCs. Also, the LSPs of the preferred embodiments of the present invention may terminate into shared or contention-based media with multiple attached MPEG devices as opposed to the point-to-point link termination of ATM PVCs or SVCs into a single MPEG device. In most cable TV networks that are the probable deployment environment of the preferred embodiments of the present invention, the LSPs may terminate at broadcast or multi-point media capable of being connected to many MPEG-capable subscriber devices.

FIG. 10 shows an embodiment of the present invention. In FIG. 10 an MPEG transport packet 1000 is encapsulated directly inside a label switched payload. The packet in FIG. 10 further comprises a header with layer 2 addressing 1002 and a label header 1011. FIG. 11 shows another embodiment wherein an MPEG transport packet 1100 is encapsulated directly inside a label switched payload found inside an ATM or frame relay header that incorporates a label within the layer 2 header 1102. One major difference in FIG. 11 from prior implementations is that the LSPs of the preferred embodiments of the present invention may be established by other protocols than the ATM layer 3 connection signaling protocol of Q.2931 that sets up ATM switched virtual circuits to carry ATM cells with MPEG payloads. These protocols for establishing LSPs and distributing labels will be discussed in more detail below. Also, FIGs. 10 and 11 are not meant to suggest that one and only one MPEG transport packet may be encapsulated into one label switched frame. The embodiments of the present invention could also work, for example, and without limitation, if multiple MPEG transport packets and/or

fractions of one MPEG transport packet are encapsulated into a single label switched frame.

Implementing Label Switching

FIG. 12 shows the common functions of standard routers. These functions can be broken down into a forwarding component 1201 and a control component 1202. The forwarding component 1201 generally contains rules for deciding how to forward packets 1211. In general, these rules specify what information in an incoming packet should be evaluated and compared to a forwarding table to decide how to forward an incoming packet. Unfortunately, in routers these rules became more and more complex to handle different types of packets and forwarding models. Some examples of forwarding models and rules for IP routers include: forwarding rules for unicast 1215, forwarding rules for multicast 1216, and forwarding rules for unicast with type of service (TOS) flags 1217. As an example, the common forwarding rule for unicast IP packets generally involves a longest match algorithm that matches the destination address of incoming packets with the entry in the forwarding table that has the longest bit length match of the destination address.

The control component 1202 of routers normally maintains the forwarding table used by the forwarding component 1201. The maintenance of this table often requires a protocol to distribute consistent information about the network among multiple routers. Thus, the control component normally includes procedures for maintaining and updating the forwarding table 1222. Examples of some common routing protocols that distribute information on routing or forwarding tables include RIP (Routing Information Protocol) 1225, OSPF (Open Shortest Path First) 1226, BGP (Border Gateway Protocol) 1227, and PIM (Protocol Independent Multicast) 1228. Those skilled in the art will realize that these examples are only a small number of the routing protocols that exist.

FIG. 13 shows the functions of a label switching router (LSR), which is used in label switching networks instead of or in addition to standard routers. In an LSR, the forwarding component 1301 and the control component 1302 have been separated into two distinct functions. This separation of functions allows simplification of the forwarding component 1301 so that it may more easily be implemented in a fast, all-hardware solution. In the forwarding component 1301, the rules for deciding how to forward packets 1311 have been reduced to one simple rule, exactly match an incoming

label to a label switching forwarding table 1315. This rule can often be implemented in a simple, indexed lookup table. Thus, the forwarding table is simplified. In an LSR the control component 1302 must still have procedures for maintaining and updating the forwarding table 1322. These procedures can be broken down into three major tasks: path determination 1325, label assignment 1335, and label distribution.

Path determination normally involves performing a mapping from a forwarding equivalence class (FEC) to a next hop (*i.e.*, a next router or LSR where the packet is forwarded to). For the IP protocol a network layer routing protocol such as RIP, OSPF, BGP, or PIM usually determines the next hop. For example, a simple non-label switched network might use RIP version 2 (RIPv2) as a routing protocol. RIPv2 is well known in the art and is used for routing in simple for IPv4 networks. The following description of RIPv2 covers the usual function of RIPv2 in routed IP networks that may not be using label switching even though the description of the function of RIPv2 is explained by terminology such as forwarding equivalence class (FEC) that is used in label switching networks. RIPv2 packets carry a list of IP network addresses identified by a 32-bit address and a 32-bit mask. In addition, in a RIPv2 packet each IP network address in the list is associated with a next hop IP address and a hop count. The IP network addresses identified by a 32-bit address and a 32-bit mask are basically a forwarding equivalence class (FEC) such that all IP packets whose destination address field falls within the range specified by an IP network address are forwarded equivalently. An IP packet that matches a forwarding equivalence class (FEC) is forwarded to a next hop router associated with that FEC. The IP address of the next hop router is specified as the next hop IP address in the RIPv2 packet. The hop count specifies the number of routers that must be transversed before a packet forwarded to the next hop router will reach the IP network specified by the FEC.

As the evolution of networks, especially from the development of fiber optics, has increased transport bandwidth relative to switching bandwidth, it becomes less expensive to carry all traffic and let the destination select the information it wants to receive as opposed to using switching technology to control the information allowed on the transport network. This type of network is similar to broadcast environments on satellite or cable networks, where a signal is sent to all or most users, but access to the signal is conditional. This type of network is much more deterministic than IP networks. Usually centralized entities determine what content is put on the network and who may have access to the content. For a system where most of the content is transported across the

network to the destination and access to the content is conditional, path determination is unlike the path discovery process used in routing protocols such as for IP. Instead the centralized entity must know the topology of the network and enable the proper paths. As part of the preferred embodiments of the present invention, a centralized system might determine paths during the label assignment process.

The other two control functions generally used in a label switching environment are label assignment, which includes procedures for creating bindings between labels and FECs, and label distribution, which includes procedures for distributing label binding information. Together these two functions work to establish and distribute FEC to label mappings. In many label switching technologies including some of the proposals for MPLS, explicitly routed label switched paths (LSPs) are envisioned for policy routing and traffic engineering. The preferred embodiments of the present invention propose to further use LSPs to set up broadcast and personal television paths for MPEG transport. To implement such a system for MPEG delivery, at least one centralized controller can accomplish the function of creating paths and assigning labels. An example of a centralized controller that could be utilized to implement such a system is the Digital Network Control System's (DNCS) Session and Resource Manager (SRM) as used in Scientific-Atlanta's Digital Broadcast Delivery System (DBDS). The DNCS SRM may implement the functions of an MPEG Digital Storage Media-Command and Control Session and Resource Manager (DSM-CC SRM), which among other things manages and controls MPEG sessions.

All references in this application to one or at least one centralized controller are not intended to be limiting as those skilled in the art will be aware that there may be requirements for multiple centralized controllers for redundancy and/or scaling reasons. In addition, those skilled in the art will be aware that the functions of a central controller may be distributed among multiple processors in multiple network devices to achieve similar or equivalent functionality.

The at least one centralized controller could base label assignments on information associated with MPEG transport packets whose format is described below. For instance, an entire MPEG transport stream might need to transverse the same label switched path (LSP), be bound for the same destination, and have identical quality of service (QoS) requirements. In such a situation all the MPEG programs in the transport stream are part of the same forwarding equivalence class (FEC). This FEC would be given an MPLS label that would designate the entire transport stream including more than

one MPEG program and the associated program specific information (PSI). This method of label assignment would be useful for transporting MPEG streams over a channel that may contain more than one MPEG transport stream because the MPEG packet IDs (PIDs) in MPEG transport packets are only unique over one MPEG transport stream. MPLS or other label switching technologies could be used to allow MPEG transport packets with identical PIDs but belonging to different MPEG transport streams to be time-division multiplexed into the same channel or medium. With MPLS or other label switching technology, multiple MPEG transport streams could be multiplexed into a channel so that a first label could be assigned to a first MPEG transport stream that includes a first Program Association Table (PAT) with MPEG PID = 0, and a second, different label could be assigned to a second MPEG transport stream that includes a second Program Association Table (PAT) with the same MPEG PID = 0. The Program Association Table (PAT) is only one example of two MPLS or other label switching forwarding equivalence classes (FECs) carrying two different MPEG transport packets that might have the same PID number. The use of PID = 0 in the preceding example is not meant to be a limitation on the MPEG PIDs that may be used in the two different MPEG transport streams carried in two different forwarding equivalence classes of a label switching network. Thus, the advantage of using MPLS or other label switching technology to carry two or more MPEG transport streams will work to differentiate MPEG transport packets belonging to different MPEG transport streams even though the MPEG transport packets have the same PID values. These same PID values for MPEG transport packets belonging to different MPEG transport streams may be any PID value including the PID = 0 of a Program Association Table (PAT) as well as any other PID value.

If a label switched path (LSP) extends all the way to a subscriber device, MPLS labeled transport streams could group the same programs that are normally provided in a frequency-division multiplexed (FDM) channel in cable TV systems. This in effect enables an evolution from the current broadband, FDM cable TV system to a full baseband, transport cable TV system. In such a network, an MPLS preselector or switch might replace the RF tuner found in current subscriber devices for FDM cable TV systems. This would allow existing MPEG transport engines to operate with baseband streams capable of much higher data rates than individual FDM channels.

Also, an MPLS label could be applied to one MPEG program that is part of an MPEG transport stream. This in effect creates a one-program MPEG transport stream and requires program specific information (PSI) for each program. This type of

assignment might be useful for unicast cases where individual subscribers desire stream control so that only one program falls in a forwarding equivalence class (FEC).

Furthermore, MPEG transport packet IDs (PIDs) identify particular flows of information such as a video program stream. These PIDs could be used by and/or assigned by at least one centralized controller as part of the label assignment function. In addition, the stream ID found in the header of an MPEG packetized elementary stream (PES) also might be used by at least one central controller as part of the label assignment process. In contrast to typical IP routing that is based on destination IP addresses, these identifiers (*i.e.*, the PID and the stream ID) in MPEG packets are more closely related to the source or content of a stream of MPEG data.

These label assignment functions available in MPLS or other similarly capable label switching technologies simplify merging streams of information that were previously dissimilar. MPEG transport already allows merging two MPEG streams to create a new MPEG stream. This type of merge might be done to merge an MPEG stream of normal program content with an MPEG stream of advertising content. MPEG transport also allows merging MPEG streams with private data streams that might contain IP datagrams. With MPLS or other similarly capable label switching technology, additional multiplexing and merging techniques are possible. For example, multiple MPEG program streams with different PIDs could be merged into one MPLS forwarding equivalence class (FEC) that contains two MPEG programs, which share the same MPLS label. In addition, it would be possible to use a single label to merge an MPEG stream with an IP datagram stream as one FEC with one label. Either the MPEG stream or the IP stream may contain normal program content while either the IP stream or the MPEG stream, respectively, may contain the advertising content. Also, it would be possible to use separate labels for the advertising stream and the program stream and have the subscriber device receive both streams. The advertising stream and the program stream may be either MPEG or IP. These previous multiplexing examples are not intended to limit the scope of the present invention in any way. Given the present disclosure, those skilled in the art will be aware that many possible ways exist for combining and multiplexing MPEG streams using label switching technologies such as MPLS.

The last main control function needed in label switching environments is label distribution. Label distribution involves informing label switching routers (LSRs) of label assignments by distributing the label binding information over communication paths. MPLS does not specify a required protocol for distributing labels so there are

various proposals for label distribution. One proposed method is to distribute label assignment information with the routing protocols by including the information in the packets of RIP, OSPF, BGP, etc. Another proposed method is to use the Resource ReSerVation Protocol (RSVP) to distribute labels. RSVP was primarily designed for informing network equipment about information flows in an IP network. Once informed about information flows by RSVP, the network equipment may allocate resources to information flows in order to implement various quality of service (QoS) criteria for an information flow. The use of RSVP as a label distribution protocol might allow label switched paths (LSPs) to be established in a network based on QoS criteria. RSVP is basically a signaling protocol for establishing QoS for integrated services over IP-oriented networks. Also, the IETF has developed a new protocol for label distribution called the Label Distribution Protocol (LDP) that is more generalized and less IP-centric than RSVP.

Constraint-based routing is one way to provide quality of service in a label switched network by determining suitable routes with the network resources to meet a variety of constraints such as a minimum bandwidth. The Constraint-based Routing Label Distribution Protocol (CR-LDP) is an extended version of LDP that enables constraint-based routing and QoS reservation in label switching networks such as an MPLS network.

The preferred embodiments of the present invention propose to use the MPEG Digital Storage Media-Command and Control (DSM-CC) protocol for label distribution to MPEG-aware devices such as MPEG-aware LSRs. Label distribution using DSM-CC may occur as part of the process of setting up DSM-CC sessions. DSM-CC is an MPEG protocol specification for signaling MPEG sessions. DSM-CC consists of two main protocols: 1) a DSM-CC user-network signaling protocol and 2) a DSM-CC user-user signaling protocol. The DSM-CC user-network protocol is designed to signal the network to setup MPEG sessions for MPEG user devices. MPEG user devices include subscriber terminal equipment, which might initiate the download of an MPEG movie. Also, MPEG user equipment includes MPEG video or audio servers, which may initiate the download of MPEG information to subscriber devices. The MPEG DSM-CC user-network protocol is somewhat similar to common layer 3 signaling protocols such as Q.931 used for circuit-switched connection establishment in narrowband ISDN or Q.2931 used for switched virtual-circuit establishment in ATM. Q.2931 is basically an extended

version of Q.931 to provide for more sophisticated quality-of-service requirements than are possible in Q.931.

However, DSM-CC was not really designed to initiate network setups. Instead, if MPEG devices are connected to a circuit switched network, then Q.931 might be used to establish circuit-switched connections between MPEG devices. Next, DSM-CC user-network signaling would establish the MPEG video or audio sessions between MPEG user devices over the circuit-switched paths. Alternatively, if MPEG devices are connected to an ATM packet-switched, virtual-circuit network, then Q.2931 might be used to establish SVCs between MPEG devices. Next, DSM-CC user-network signaling would establish the MPEG video or audio sessions between MPEG user devices over the SVCs. The DSM-CC user-user protocol is designed to send control information between MPEG user devices. Thus, the DSM-CC user-user protocol may used by an MPEG subscriber device to control a video flow from an MPEG video source. The subscriber device may be able to use VCR-like functional control to pause, rewind, or fast-forward video from an MPEG video source using the DSM-CC user-user protocol.

The DSM-CC standard covers various types of sessions such as Continuous Feed Sessions (CFS) and Exclusive Sessions (ES). In addition, the DSM-CC standard covers user to network (UN) messages that generally create a user-network protocol as well as user to user (UU) messages that generally create a user-user protocol. Presently the DSM-CC standard does not define network to network messages. If an external server wants to set up a continuous feed session (CFS), then it uses one or more user to network (UN) messages. Similarly, if an external client wants to create an exclusive session (ES), then it also uses one or more user to network (UN) messages. If a client wants to control a server, then one or more user to user (UU) messages are communicated.

When establishing DSM-CC sessions, allocation of resources in the network may be accomplished through the communication of other protocols besides DSM-CC. For example, if the network wishes to set up a continuous feed session (CFS) between an MPEG client and an MPEG server, then various non-DSM-CC messages may be used to establish the CFS. The protocols carrying these messages might include remote procedure call (RPC) and/or simple network management protocol (SNMP) as examples of protocols that may be used to communicate resource allocations in a CATV network between equipment such as a controller and a modulator. These additional protocol messages might also be used to distribute labels during the process of setting up and managing DSM-CC sessions. Those skilled in the art will be aware that RPC and SNMP

are just examples of possible protocols that might be used in the label distribution, and the embodiments of the present invention are not intended to be limited to those protocols. Thus, the labels may be distributed in DSM-CC messages or other non-DSM-CC protocols as part of the process of establishing DSM-CC sessions.

Furthermore, other label distribution protocols such as LDP, CR-LDP, and/or RSVP could be used to distribute labels to non-MPEG-aware network equipment (*e.g.*, LSRs) in a label switching network. Although those skilled in the art will be aware that other protocols also may be used, the preferred embodiments of the present invention propose that the label distribution messages for MPEG DSM-CC and other protocols for label distribution may be carried in IP datagrams because of the generally ubiquitous nature of basic IP functionality in network devices. Basic IP functionality is often included in network devices to allow such common capabilities as network management. Still, the embodiments of the present invention are in no way limited to only distributing label information in IP datagrams.

In the embodiments of the present invention, standard IP messages may be used to implement the MPLS control component functions of path or route determination, label assignment, and label distribution. Thus, the devices and methods of the embodiments of the present invention may be used in a network that carries IP network level traffic. However, the embodiments of the present invention make it possible to carry MPEG transport packets directly within MPLS or other label switched technology without the necessity of placing the MPEG transport packets inside IP datagrams or other network level protocol messages.

MPEG Transport Packets

MPEG, the Motion Pictures Expert Group, has established standard procedures for encoding information streams including real-time video and audio. Furthermore, some of these standards include protocols for the transport and signaling of information flows. This present application has been written mainly based on some of the MPEG-2 standards. However, this is not meant to indicate any limitations on the present invention. Thus, it is expected that the preferred embodiments of the present invention will work without limitation with MPEG-1, MPEG-2, MPEG-4, MPEG-7, and MPEG-21 as well as MPEG and non-MPEG standards to handle real-time information flows that may be developed in the future.

MPEG-2 includes a Systems Specification 1401 that defines some of the interfaces for MPEG packets as shown in FIG. 14. FIG. 14 shows one interface of the MPEG Systems Specification 1401 as the dashed line 1405. Video data 1415 goes through video encoder 1417 into the MPEG systems specification 1401. This MPEG encoded video from video encoder 1417 is output on connection 1419 and is input into packetizer 1421 to create a video packetized elementary stream (PES) 1425. In addition, FIG. 14 shows audio data 1435 going through audio encoder 1437 into MPEG systems specification 1401. This MPEG encoded audio from audio encoder 1437 is output on connection 1439 and is input into packetizer 1441 to create audio packetized elementary stream (PES) 1445. Video PES 1425 is combined with audio PES 1445 in MPEG program stream mux (or multiplexer) 1455 to create MPEG program stream 1465. Furthermore, video PES 1425 is combined with audio PES 1445 in MPEG transport stream mux (or multiplexer) 1475 to create MPEG transport stream 1485.

An MPEG elementary stream is an output from an audio or video encoder. Before communicating this MPEG elementary stream, it must be packetized into a packetized elementary stream (PES). MPEG systems specification 1401 comprises rules for creating MPEG program streams 1465 and MPEG transport streams 1485. MPEG program streams 1465 include information to carry a single MPEG program. An MPEG program comprises different PES streams that share a common time base. For example, an MPEG movie may include a video PES, an English audio track, and a Spanish audio track all as one program. The MPEG program stream 1465 from MPEG program stream multiplexer 1455 is designed to allow local delivery of MPEG programs to players such as a VCR-like unit playing a movie on a locally connected TV. The MPEG transport stream 1485 from MPEG transport stream multiplexer 1475 is designed to allow delivery of MPEG programs across communication networks in MPEG transport packets. The MPEG transport protocol is designed to carry one or more MPEG programs and to communicate the programs in MPEG transport packets. In addition to carrying packetized elementary streams (PES), the MPEG program and the MPEG transport specifications are designed to carry additional information such as private data, time synchronization information, and service and control information. An example of private data might be closed captioning for an MPEG movie. The private data portion of MPEG transport may even carry IP datagrams.

FIG. 15 shows how elementary streams of video are placed in packetized elementary streams (PES) and MPEG transport packets. Elementary stream 1505 is an I-

Picture, which is an encoded "intra" frame of video that is part of the algorithm used in MPEG video compression. Elementary stream 1506 is a P-Picture, which is an encoded "predictive" frame of video that is part of the algorithm used in MPEG video compression. Nothing in this application is intended to limit the embodiments of the present invention only to methods of digitizing and compressing audio or video data that are currently described in MPEG specifications. Thus, other digitizing and/or compression algorithms than those in MPEG specifications might also be used, and the MPEG I-Picture/P-Picture compression algorithm is only an example. In FIG. 15 the I-Picture elementary stream 1505 is packetized into a packetized elementary stream (PES) comprising PES Header 1515 and PES Data 1516, which includes the data from the I-picture frame 1505. The P-Picture elementary stream 1506 is packetized into a packetized elementary stream (PES) comprising PES Header 1525 and PES Data 1526, which includes the data from the P-picture frame 1506. The packets of a PES include a stream ID in the header of PES packets. These stream IDs allow determination of the type of information contained in a PES packet.

FIG. 15 further shows how one variable length PES packet comprising PES header 1515 and PES packet payload 1516 and containing data from an I-picture elementary stream is encapsulated into several fixed length MPEG transport stream packets. MPEG transport stream packet with transport stream header 1535 contains an MPEG transport stream packet payload 1536 encapsulating PES header 1515 and the first part of the packetized data from the I-Picture 1516. MPEG transport stream packet with transport stream header 1545 contains an MPEG transport stream packet payload 1546 encapsulating the second part of the packetized data from the I-Picture 1516. Finally, MPEG transport stream packet with transport stream header 1555 contains an MPEG transport stream packet payload 1556 encapsulating the final part of the packetized data from the I-Picture 1516 as well as stuffing bits to fill out the MPEG transport packet to its fixed length size.

The MPEG transport stream packet header includes a 13 bit field known as a packet ID or PID that is used to identify information in the MPEG transport multiplexing hierarchy. This PID field can include values from 0 to 2^{13} - 1 or 8191 (decimal). Some of the PID values such as 0 and 1 are assigned to specific tables. PID values 2 to 15 (decimal) are reserved for future use. PID value 8191 (decimal) designates a null packet, while PID values 16 to 8190 (decimal) are available for PES streams, map tables, and network tables.

FIG. 16 shows an example of how the multiplexing architecture of MPEG transport works. MPEG transport packets with PID = 0 carry a Program Association Table (PAT) 1605 that contains a list of MPEG programs associated with other PID values. MPEG transport packets with PID = 1 carry a Conditional Access Table 1615 that is used to limit access to restricted and encrypted programs such as in a pay-per-view service. The first entry in the Program Association Table 1605 associates program 0 with PID 20. This means that MPEG transport packets with PID = 20 contain a Network Information Table (NIT) 1625. The Network Information Table (NIT) 1625 contains information about the network properties of the network carrying the MPEG transport packets. The Network Information Table (NIT) 1625 is generally associated with program 0 in the Program Association Table (PAT) 1605. Program Association Table (PAT) 1605 further shows an example of three MPEG programs: program 1, program 22, and program 35.

From the Program Association Table (PAT) 1605, MPEG multiplexers/ demultiplexers can determine that for MPEG program 1 the Program Map Table (PMT) 1635 may be found in MPEG transport packets with PID = 68. As shown in FIG. 16, Program Map Table 1635 for program 1 contains a table listing the Packetized Elementary Streams (PES) included in program 1. For Program Map Table 1635 the first PES stream is a video stream and can be found in MPEG transport packets with PID = 190 as shown as PES 1637. In addition, for Program Map Table 1635 the second PES stream is an audio stream and can be found in MPEG transport packets with PID = 417 as shown as PES 1639.

Also, from the Program Association Table (PAT) 1605, MPEG multiplexers/ demultiplexers can determine that for MPEG program 22 the Program Map Table (PMT) 1645 may be found in MPEG transport packets with PID = 150. As shown in FIG. 16, Program Map Table 1645 for program 22 contains a table listing the Packetized Elementary Streams (PES) included in program 22. For Program Map Table 1645 the first PES stream is a video stream and can be found in MPEG transport packets with PID = 75 as shown as PES 1647. In addition, for Program Map Table 1645 the second PES stream is an audio stream and can be found in MPEG transport packets with PID = 235 as shown as PES 1649. Finally, for Program Map Table 1645 the third PES stream is a private data stream and can be found in MPEG transport packets with PID = 512 as shown as PES 1651.

Furthermore, from the Program Association Table (PAT) 1605, MPEG multiplexers/demultiplexers can determine that for MPEG program 35 the Program Map Table (PMT) 1655 may be found in MPEG transport packets with PID = 184. As shown in FIG. 16, Program Map Table 1655 for program 35 contains a table listing the Packetized Elementary Streams (PES) included in program 35. For Program Map Table 1655 the first PES stream is a video stream and can be found in MPEG transport packets with PID = 34 as shown as PES 1657. In addition, for Program Map Table 1655 the second PES stream is an audio stream and can be found in MPEG transport packets with PID = 200 as shown as PES 1659. Finally, for Program Map Table 1645 the third PES stream is an audio stream and can be found in MPEG transport packets with PID = 251 as shown as PES 1661.

CATV Network

FIG. 17 shows an example of how a CATV network might utilize MPEG transport packets in a label switching header. The network of FIG. 17 is only meant to be one example of the many possible networks. Those skilled in the art will be aware of additional implementations, which are intended to be included in the present invention. Some of the equipment in the network of FIG. 17 would be MPEG-aware (*i.e.*, capable of interpreting MPEG transport packets within the label switching header) while other equipment may not be able to directly decode MPEG transport packets (*i.e.*, be non-MPEG-aware). For instance, Label Switching Router (LSR) 1705 would likely be located at a headend or a central office (CO). LSR 1705 would be MPEG-aware and take input from co-located MPEG source 1707 or co-located storage 1709 that may have stored or cached MPEG programs. In addition, other MPEG sources such as MPEG source 1711 may be connected to LSR 1705 through various networks such as long distance delivery network 1715. The long distance network 1715 may or may not be MPEG-aware. The long distance delivery network 1715 may be made up of other label switching routers, SONET rings, or a virtual private network with a guaranteed quality of service that is tunneled through the public Internet. These are just some examples of delivery networks that may be used to implement long distance delivery network 1715, and none of these examples is intended to limit the present invention in any way.

Label Switching Router (LSR) 1705 would be further connected to a local delivery network 1725. One example, without limitation, of a delivery network is described in U.S. Patent Application entitled "A System Architecture For Customized CATV Services," by Luis A. Rovira, Lorenzo Bombelli, Paul Connolly, Donald L. Sipes, Jr., and Douglas Woodhead, filed on the same day as the present application with Attorney Docket No. A-6551, and is hereby incorporated by reference in its entirety. This network would connect LSR 1705 to other MPEG-aware LSRs such as LSR 1735. In addition, the network would work with non-MPEG-aware LSRs such as LSR 1745, which is also connected to local delivery network 1725. MPEG-aware LSR 1735 could be connected in a first distribution hub to distribution interface 1755, which may be MPEG-aware and which is further connected to subscriber distribution network 1757. Subscriber device 1765 is an example of an MPEG-aware device connected to subscriber distribution network 1757. Non-MPEG-aware LSR 1745 could be connected in a second distribution hub to distribution interface 1775, which may be MPEG-aware and which is

further connected to subscriber distribution network 1777. Subscriber device 1785 is an example of an MPEG-aware device connected to subscriber distribution network 1777. Finally, MPEG-aware LSR 1705 would be connected to an MPEG Digital Storage Medium-Command and Control (DSM-CC) Session and Resource Management (SRM) server 1795, which would also likely be co-located at a headend or central office (CO).

The DSM-CC SRM server 1795 can be used for managing MPEG session establishment and session release. The DSM-CC protocol can be used to distribute labels to MPEG-aware LSRs 1705 and 1735. Other label distribution protocols such as LDP (Label Distribution Protocol), CR-LDP (Constraint-based Routing-LDP), and/or RSVP (Resource ReSerVation Protocol) can be used to distribute labels to non-MPEG-aware LSRs such as LSR 1745.

Example Applications for the Present Invention

It is expected that the preferred embodiments of the present invention will allow simpler deployment of many new services over the cable TV network. Some of these new services involve the use of subscriber-customized information. A new set of services known generally as personal TV is a non-limiting example of how subscriber-customized information might use the enhanced network capabilities resulting from the preferred embodiments of the present invention. Personal TV encompasses at least five basic functions or abilities. A first function of personal TV might be broadcast-on-demand or MPEG multicast, which involves the ability to deliver, simultaneously to multiple subscribers who request it, content that is not broadcast to all subscribers or does not occupy the broadcast portion of the CATV spectrum. Another function of personal TV might be the ability to deliver stored content on demand from a number of storage locations. This function is often called video-on-demand (VOD), but it may encompass other types of programs besides movie videos. A third function of personal TV might be the ability to retrieve broadcast content whose simultaneous broadcast time has already passed. Although this time shifting function could be implemented by a subscriber's customer premise equipment (CPE) such as a video cassette recorder (VCR) or other form of personal video recorder (PVR), providing this capability with equipment in the network allows a subscriber to watch a previously broadcast program even though the subscriber did not previously anticipate the need to record the content. A fourth function of personal TV might be the ability of subscribers to exercise control over broadcast

streams, which may include VCR-like features of pause and replay. These features could be implemented by creating a unicast stream for a subscriber or user that has initiated a pause or replay request. A fifth function of personal TV might be the ability to insert targeted or customized advertising into streams.

Such customized advertising screens may be carried in IP datagrams or in other existing or yet to be developed protocols. The ability to layer communication protocols means that the IP datagrams may be delivered to user devices through many methods. One example includes delivering the IP datagrams in packets without label switching headers. Another example may be delivering IP datagrams in packets with label switching headers. As discussed above, some LSRs dynamically allocate labels to IP datagram flows based upon the number of IP datagrams reaching a threshold count in a specified amount of time. In the case of dynamically allocated labels, the IP datagrams may or may not be carried in label switched frames. Also, even if labels are not allocated dynamically to flows of IP datagrams, the IP datagrams may or may not be carried in label switched frames. In addition, IP datagrams might be carried in the private data streams of MPEG transport packets. None of these examples of carrying IP datagrams are intended to be limiting in any way. They are only meant as examples of how IP datagrams may be multiplexed with MPEG transport in at least one embodiment of the present invention, which involves natively or directly carrying MPEG transport packets inside of label switched frames such as the frames of MPLS.

FIG. 17 will be referenced in the explanation of the following examples. Also, MPLS will be assumed to be the label switching protocol used by local delivery network 1725. However, the techniques of the present invention are not necessarily limited to MPLS. Thus, other label switching protocols capable of carrying MPEG might also be used. As an example of operation of the preferred embodiments of the present invention, assume that the headend or CO system operator has identified content from MPEG source 1707 or MPEG source 1711 that should be delivered to subscriber device 1765 on subscriber distribution network 1757. If the content is to be placed on the broadcast or narrowcast portion of distribution network 1757, then label switching of MPEG transport as described in the preferred embodiments of the present invention is not required. Instead MPEG would be broadcast as normally done for carrying multiple channels to subscriber locations such as through current FDM and TDM techniques in the cable network. However, if the content from MPEG source 1707 or MPEG source 1711 is to implement a personal TV service such as broadcast-on-demand, then the content will be

switched by and/or pass through at least some of the LSRs of local delivery network 1725. LSR 1735 is an MPEG-aware label switching router that is specifically designed for interpreting and acting on MPEG transport in MPLS. In contrast, LSR 1745 is an MPLS compliant label switching router that is neither aware of nor capable of acting on the MPEG content within the MPLS frames.

The MPEG packets may only be addressed to subscribers who request the content, but many subscribers may want to view this content in real-time, simultaneously with the subscribers that requested the broadcast-on-demand. The operator may use the DSM-CC Session and Resource Manager (SRM) 1795 to preconfigure sessions; thus establishing a path and the required resources to deliver this content to the subscribers on subscriber distribution network 1757. Such a preconfiguration would be done by establishing a DSM-CC Continuous Feed (CF) session. Subscribers on subscriber distribution network 1757 may use a DSM-CC Client Session Setup Request to connect to this service and to be bound to the existing CF session. Also, the conditional access (CA) process from the MPEG transport Conditional Access Table (CAT) may be activated to ensure proper authorization to decrypt the broadcast-on-demand content. Use of the conditional access table in a network is optional and does not have to be implemented to obtain the benefits of the embodiments of the present invention.

Suppose that subscriber device 1785 on subscriber distribution network 1777 now wants to receive the same content and issues a DSM-CC Session Setup Request. Because the content would be simultaneously delivered to other subscribers on subscriber distribution network 1777 who may request the content, a second Continuous Feed (CF) session is set up involving distribution interface 1775 and non-MPEG-aware LSR 1745. Any other subscribers on subscriber distribution network 1777 that desire this content would be bound to this second CF session. If local delivery network 1725 is implemented as a ring architecture, then MPLS traffic may be bound for more than one LSR such as LSRs 1735 and 1745. In a network of this type, an MPLS label associated with a multicast group may normally be assigned at the creation of any continuous feed session.

Furthermore, suppose that a second headend or CO (though not shown in FIG. 17) were connected to long distance delivery network 1715. The equipment co-located in this second headend might be similar to the equipment co-located at the headend shown in FIG. 17 (*i.e.*, the first headend). This equipment in the first headend possibly includes LSR 1705, MPEG source 1707, storage 1709, and DSM-CC SRM 1795. The second headend (not shown in FIG. 17) would have similar equipment. Also, in this second

headend the LSRs would be connected by another local delivery network similar to local delivery network 1725 connected to LSR 1705 as shown in FIG. 17. If equipment considered part of the assumed second headend attempted to negotiate a continuous feed session for the same multimedia stream from MPEG source 1711 that was already allocated a continuous feed session by DSM-CC SRM 1795 in the first headend, then MPEG source 1711 could provide the DSM-CC SRM in the second headend (not shown) with the same MPLS label already used for the content from MPEG source 1711 by the first headend. The DSM-CC SRM in the second headend could use the provided MPLS label or select a new label. Label conflicts between the two headends can be resolved because MPLS labels have only local significance to a particular network link. Label switching routers (LSRs) may change the labels of a packet as it comes into the LSR from one link and leaves the LSR through another link.

MPLS or other label switching protocols with similar features allows the previous example to work properly with packets passing through both MPEG-aware and non-MPEG-aware equipment. The network should both assign MPLS labels to broadcast-on-demand packets as well as distribute the labels to any equipment in the signal path. In the preferred embodiments of the present invention DSM-CC SRM 1795 may assign the MPLS label to the stream during the DSM-CC session setup process. The DSM-CC SRM 1795 may then use MPEG DSM-CC signaling to distribute the label to MPEG-aware, DSM-CC compliant equipment in the signal path. Distribution interfaces 1755 and 1775 may terminate the label switched paths (LSPs) of the label switching network or the LSPs may extend all the way to subscriber devices such as 1765 and 1785. The DSM-CC signaling messages could be carried in IP datagrams that are directly addressed to specific equipment involved in a continuous feed (CF) session. In addition, the DSM-CC SRM 1795 could be used to initiate standard signaling such as CR-LDP and to provide the information for a constrained path. A constrained path is a path for a flow of information packets that is constrained to those path segments that meet the quality of service requirements of the information flow.

As another example of implementing personal television functions using the preferred embodiments of the present invention, suppose that the headend operator has identified the content for a broadcast-on-demand continuous feed session to be a candidate for stream control by subscribers. Such stream control may include VCR-like functions of pausing, rewinding, and fast-forward searching. In order to accomplish the stream control functions, the content of the program must be cached in a device such as

storage 1709 before a subscriber chooses to rewind the program. Thus, when the continuous feed (CF) session is established, traffic designated for potential stream control by subscribers would likely be routed to storage 1709 for caching in addition to being routed to subscribers.

When a subscriber requests a stream control function such as pause or instant replay, the client application must initiate a chain of events in a subscriber device such as 1785 that results in a DSM-CC Client Session Setup Request. DSM-CC SRM 1795 would receive the request and setup a DSM-CC Exclusive Session (ES) between storage 1709 and the requesting subscriber device 1785 including any MPEG-aware components in the signal path such as distribution interface 1775. MPLS labels associated with the exclusive session (ES) could be distributed in session setup messages to MPEG-aware devices such as storage 1709, LSR 1705, and distribution interface 1775. MPLS labels associated with this exclusive session (ES) could be distributed to non-MPEG-aware devices such as LSR 1745 using the Label Distribution Protocol (LDP) or by other means. Stream control commands such as pause or rewind between subscriber device 1785 and storage 1709 may use the command set of the DSM-CC user-user signaling protocol, or they may be implemented with other commands that have been designed into the client and server applications. These example applications are only meant to suggest some of the various ways that the embodiments of the present invention may be used to advantageously allow deployment of networks with improved functionality. Those skilled in the art will be aware of other possibilities suggested by the embodiments of the present invention.

Note that the embodiments of the present invention may be implemented using at least one dedicated logical circuit, at least one processor and software, or at least one combination circuit that includes at least one processor circuit with software and at least one specific dedicated circuit. Those, skilled in the art of implementing protocol specifications will be familiar with the design tradeoffs of implementing hardware and software functions. It is understood that all such permutations of various implementations are included herein.

The software, which comprises an ordered listing of executable instructions for implementing logical functions, can be embodied in any computer-readable medium for use by or in connection with an instruction execution system, apparatus, or device, such as a computer-based system, processor-containing system, or other system that can fetch the instructions from the instruction execution system, apparatus, or device and execute

the instructions. In the context of this document, a "computer-readable medium" can be any means that can contain, store, communicate, propagate, or transport the program for use by or in connection with the instruction execution system, apparatus, or device. The computer readable medium can be, for example but not limited to, an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system, apparatus, device, or propagation medium. More specific examples (a nonexhaustive list) of the computer-readable medium would include the following: an electrical connection (electronic) having one or more wires, a portable computer diskette (magnetic), a random access memory (RAM) (electronic), a read-only memory (ROM) (electronic), an erasable programmable read-only memory (EPROM or Flash memory) (electronic), an optical fiber (optical), and a portable compact disc read-only memory (CDROM) (optical). Note that the computer-readable medium could even be paper or another suitable medium upon which the program is printed, as the program can be electronically captured, via for instance optical scanning of the paper or other medium, then compiled, interpreted or otherwise processed in a suitable manner if necessary, and then stored in a computer memory.

CLAIMS

Now, therefore, at least the following is claimed:

1. A method of encapsulating MPEG packets within a label switching protocol, the method comprising the steps of:
   implementing a label switching protocol layer, the label switching protocol layer allowing a plurality of label switching devices to forward label switched frames based upon a label field in the label switched frames;
   implementing an MPEG protocol layer, the MPEG protocol layer capable of carrying multimedia information; and
   encapsulating the MPEG protocol layer inside the label switching protocol layer.

2. The method of claim 1, wherein the MPEG protocol layer is an MPEG transport layer.

3. The method of claim 1, wherein the MPEG protocol layer is encapsulated directly inside the label switching protocol layer.

4. The method of claim 1, wherein the label switching protocol layer is a multiple protocol label switching layer.

5. The method of claim 4, wherein the multiple protocol label switching layer is capable of being encapsulated by multiple protocols.

6. The method of claim 4, wherein the multiple protocol label switching layer is capable of encapsulating multiple protocols.

7. The method of claim 6, wherein the multiple protocol label switching layer is used to encapsulate both the MPEG protocol layer as well as a second protocol layer using a label that has the same value for both encapsulations.

8. The method of claim 7, wherein the MPEG protocol layer comprises program content while the second protocol layer comprises advertising content.

9. The method of claim 7, wherein the MPEG protocol layer comprises advertising content while the second protocol layer comprises programs content.

10. The method of claim 7, wherein the second protocol layer is an Internet Protocol (IP) layer.

11. The method of claim 4, wherein the multiple protocol label switching layer is the Multi-Protocol Label Switching (MPLS) protocol that is defined by the Internet Engineering Task Force (IETF).

12. The method of claim 1, wherein a label is a value that can be assigned to the label field of the label switched frames, the label being distributed to the plurality of label switching devices using at least one protocol for label distribution.

13. The method of claim 12, wherein the plurality of label switching devices comprises at least one MPEG-aware label switching device, the at least one protocol for label distribution comprising a protocol for establishing and releasing sessions between MPEG equipment, the protocol for establishing and releasing sessions between MPEG equipment being used to distribute the label to the at least one MPEG-aware label switching device.

14. The method of claim 13, wherein the protocol for establishing and releasing sessions between MPEG equipment is a Digital Storage Media-Command and Control (DSM-CC) user-network protocol as defined by the Motion Pictures Expert Group (MPEG).

15. The method of claim 13, wherein the plurality of label switching devices further comprises at least one non-MPEG-aware label switching device, the label being distributed to the at least one non-MPEG-aware label switching devices using a different protocol for label distribution than the protocol for establishing sessions between MPEG equipment.

16. The method of claim 1, wherein the label switching protocol layer is capable of encapsulating a first MPEG transport stream and a second MPEG transport stream in the same communications medium by using a first label for the first MPEG transport stream and a second label for the second MPEG transport stream, the second label being different from the first label, the first MPEG transport stream using Packet IDs (PIDs) that are unique within the first MPEG transport stream, and the second MPEG transport stream using Packet IDs (PIDs) that are unique within the second MPEG transport stream.

17. A device for communicating MPEG packets in a label switched environment, the device comprising:

    logic configured to implement a label switching protocol layer, the label switching protocol layer allowing a plurality of label switching devices to forward label switched frames based upon a label field in the label switched frames;

    logic configured to implement an MPEG protocol layer, the MPEG protocol layer capable of carrying multimedia information; and

    logic configured to encapsulate the MPEG protocol layer inside the label switching protocol layer.

18. The device of claim 17, wherein the MPEG protocol layer is an MPEG transport layer.

19. The device of claim 17, wherein the MPEG protocol layer is encapsulated directly inside the label switching protocol layer.

20. The device of claim 17, wherein the label switching protocol layer is a multiple protocol label switching layer.

21. The device of claim 20, wherein the multiple protocol label switching layer is capable of being encapsulated by multiple protocols.

22. The device of claim 20, wherein the multiple protocol label switching layer is capable of encapsulating multiple protocols.

23. The device of claim 22, wherein the multiple protocol label switching layer is used to encapsulate both the MPEG protocol layer as well as a second protocol layer using a label that has the same value for both encapsulations.

24. The device of claim 23, wherein the MPEG protocol layer comprises program content while the second protocol layer comprises advertising content.

25. The device of claim 23, wherein the MPEG protocol layer comprises advertising content while the second protocol layer comprises programs content.

26. The device of claim 23, wherein the second protocol layer is an Internet Protocol (IP) layer.

27. The device of claim 20, wherein the multiple protocol label switching layer is the Multi-Protocol Label Switching (MPLS) protocol that is defined by the Internet Engineering Task Force (IETF).

28. The device of claim 17, wherein a label is a value that can be assigned to the label field of the label switched frames, the label being distributed to the plurality of label switching devices using at least one protocol for label distribution.

29. The device of claim 28, wherein the plurality of label switching devices comprises at least one MPEG-aware label switching device, the at least one protocol for label distribution comprising a protocol for establishing and releasing sessions between MPEG equipment, the protocol for establishing and releasing sessions between MPEG equipment being used to distribute the label to the at least one MPEG-aware label switching device.

30. The device of claim 29, wherein the protocol for establishing and releasing sessions between MPEG equipment is a Digital Storage Media-Command and Control (DSM-CC) user-network protocol as defined by the Motion Pictures Expert Group (MPEG).

31. The device of claim 29, wherein the plurality of label switching devices further comprises at least one non-MPEG-aware label switching device, the label being distributed to the at least one non-MPEG-aware label switching devices using a different protocol for label distribution than the protocol for establishing sessions between MPEG equipment.

32. The device of claim 17, wherein the label switching protocol layer is capable of encapsulating a first MPEG transport stream and a second MPEG transport stream in the same communications medium by using a first label for the first MPEG transport stream and a second label for the second MPEG transport stream, the second label being different from the first label, the first MPEG transport stream using Packet IDs (PIDs) that are unique within the first MPEG transport stream, and the second MPEG transport stream using Packet IDs (PIDs) that are unique within the second MPEG transport stream.

33. A method of communicating a packetized stream of digital multimedia information in at least one label switched frame that is communicated through a network, the method comprising the steps of:

obtaining a packetized stream of digital multimedia information, the packetized stream of digital multimedia information comprising at least one packet that represents at least a portion of the packetized stream of digital multimedia information;

placing the at least one packet of the packetized stream of digital multimedia information inside at least one label switched frame; and

transmitting the at least one label switched frame into the network.

34. The method of claim 33, wherein the obtaining step comprises the step of generating the packetized stream of digital multimedia information.

35. The method of claim 33, wherein the obtaining step comprises the step of receiving the packetized stream of digital multimedia information from at least one input source.

36. The method of claim 33, wherein the at least one packet of the packetized stream of digital multimedia information conforms to an MPEG transport protocol.

37. The method of claim 36, wherein the at least one label switched frame conforms to a Multi-Protocol Label Switching (MPLS) protocol.

38. The method of claim 37, wherein the at least one packet of the packetized stream of digital multimedia information is placed directly inside the at least one label switched frame so that the MPEG transport protocol is natively carried inside the MPLS protocol.

39. A method of recovering a packetized stream of digital multimedia information from at least one label switched frame that is communicated through a network, the method comprising the steps of:

receiving at least one label switched frame from the network;

retrieving at least one packet that represents at least a portion of the packetized stream of digital multimedia information from within the at least one label switched frame; and

rebuilding parts of the packetized stream of digital multimedia information based upon the least one packet retrieved from within the at least one label switched frame.

40. The method of claim 39, wherein the at least one packet of the packetized stream of digital multimedia information conforms to an MPEG transport protocol.

41. The method of claim 40, wherein the at least one label switched frame conforms to a Multi-Protocol Label Switching (MPLS) protocol.

42. The method of claim 41, wherein the at least one packet of the packetized stream of digital multimedia information is retrieved from directly inside the at least one label switched frame so that the MPEG transport protocol is natively carried inside the MPLS protocol.

43. A communication signal, embodied in a computer-readable medium, comprising:

    at least one packet of a packetized stream of digital multimedia information; and

    at least one label switched frame that is used to encapsulate the at least one packet of the packetized stream of digital multimedia information, the at least one label switched frame comprising information for forwarding the communication signal along a label switched path (LSP) in a network capable of propagating the communication signal.

44. The communication signal of claim 43, wherein the at least one packet of the packetized stream of digital multimedia information conforms to an MPEG transport protocol.

45. The communication signal of claim 44, wherein the at least one label switched frame conforms to a Multi-Protocol Label Switching (MPLS) protocol.

46. The communication signal of claim 45, wherein the at least one packet of the packetized stream of digital multimedia information is encapsulated directly inside the at least one label switched frame so that the MPEG transport protocol is natively carried inside the MPLS protocol.

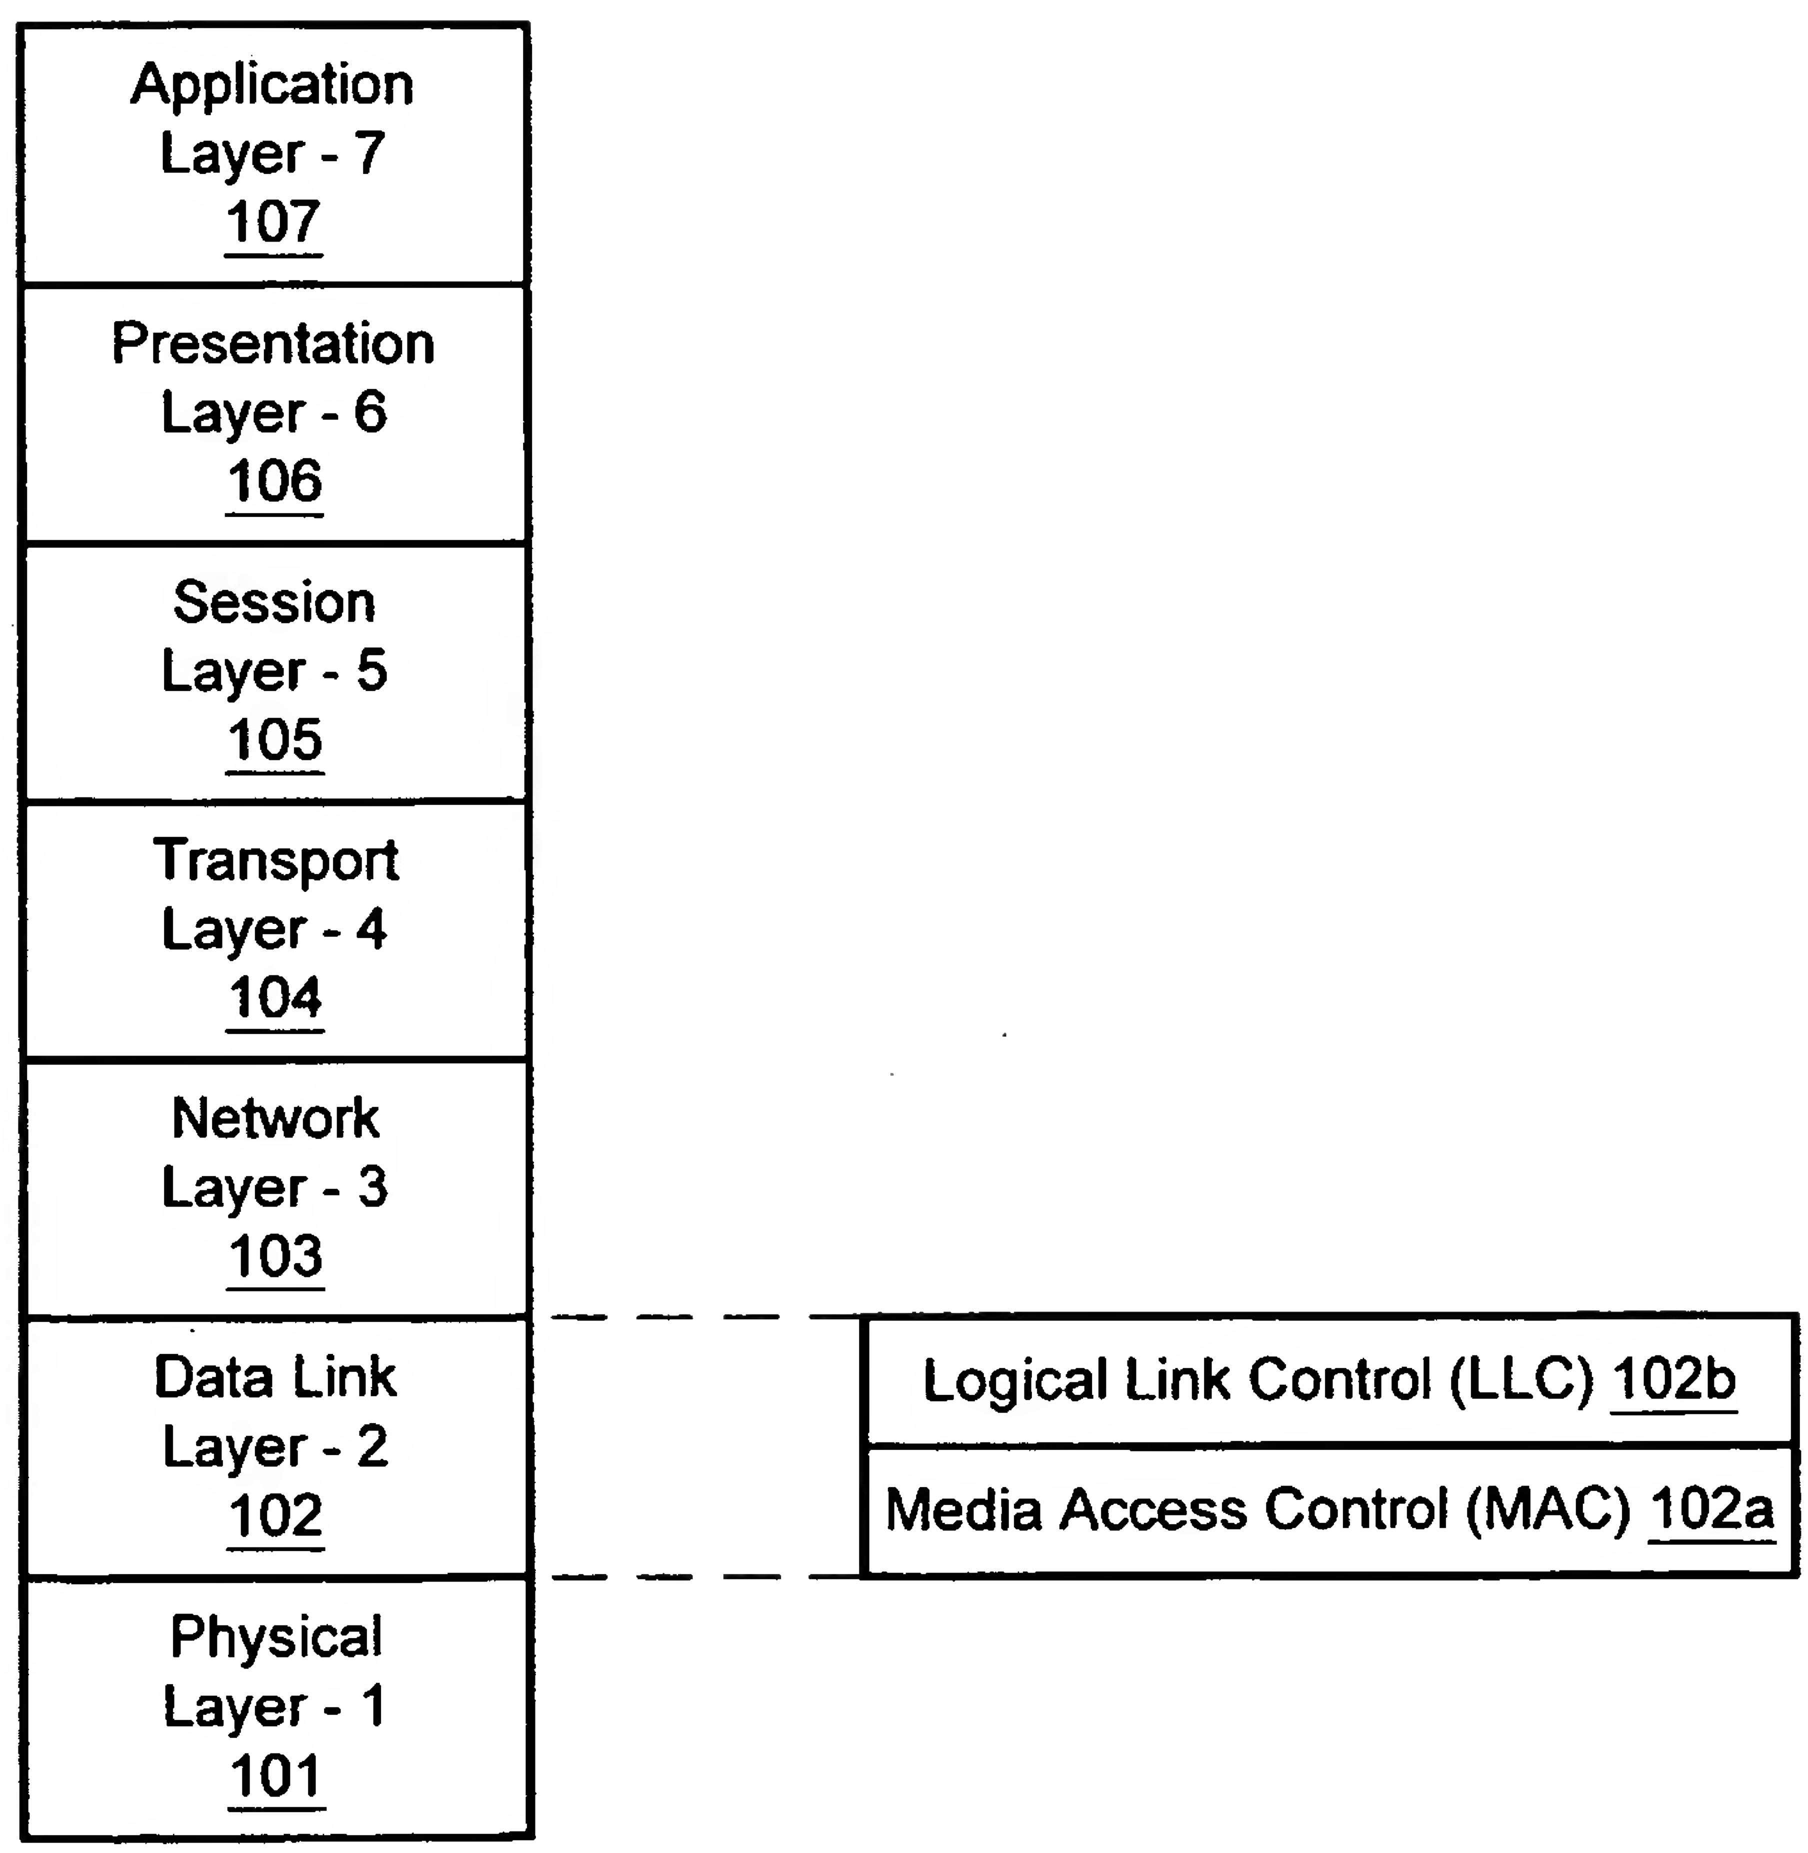

FIG. 1 - Prior Art

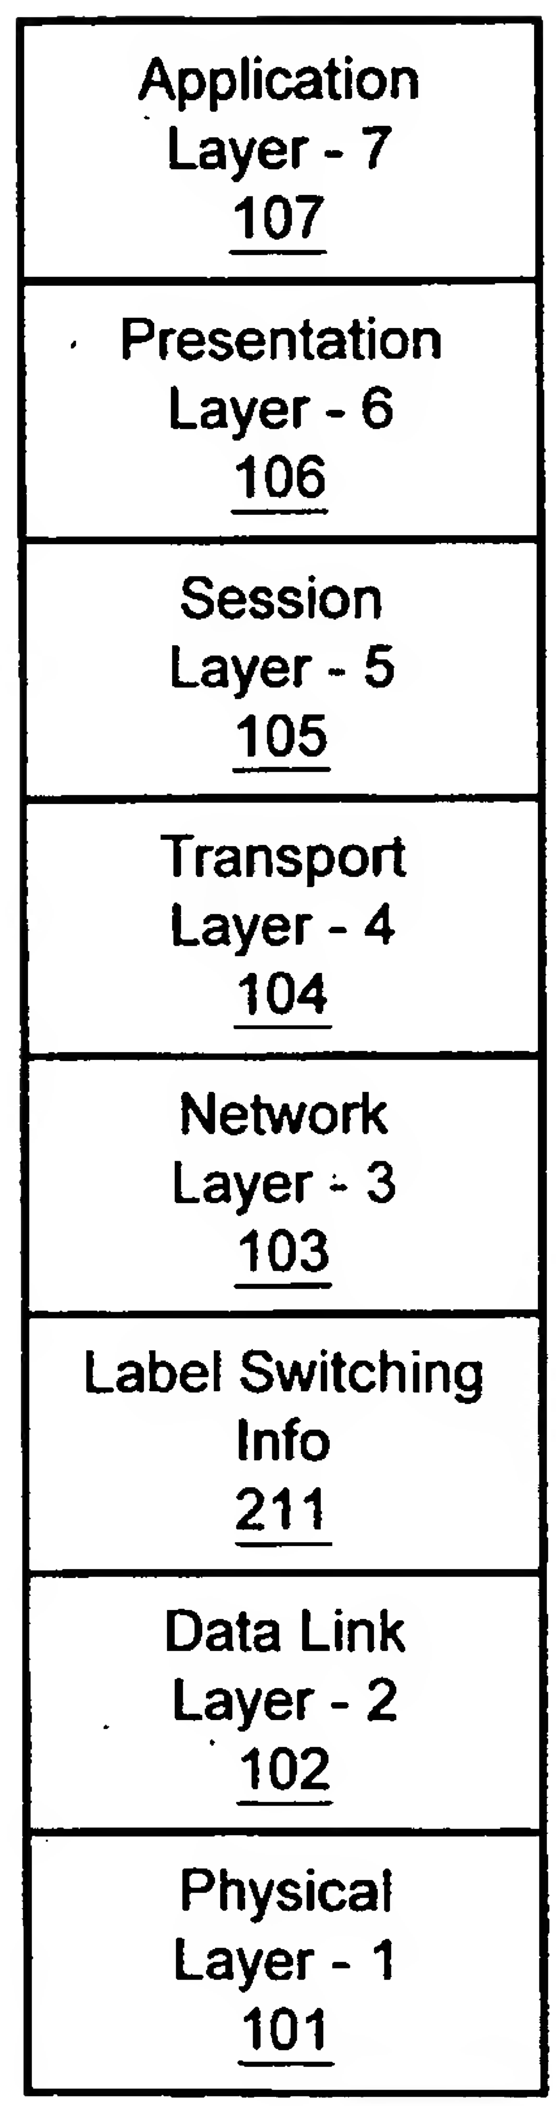

FIG. 2 - Prior Art

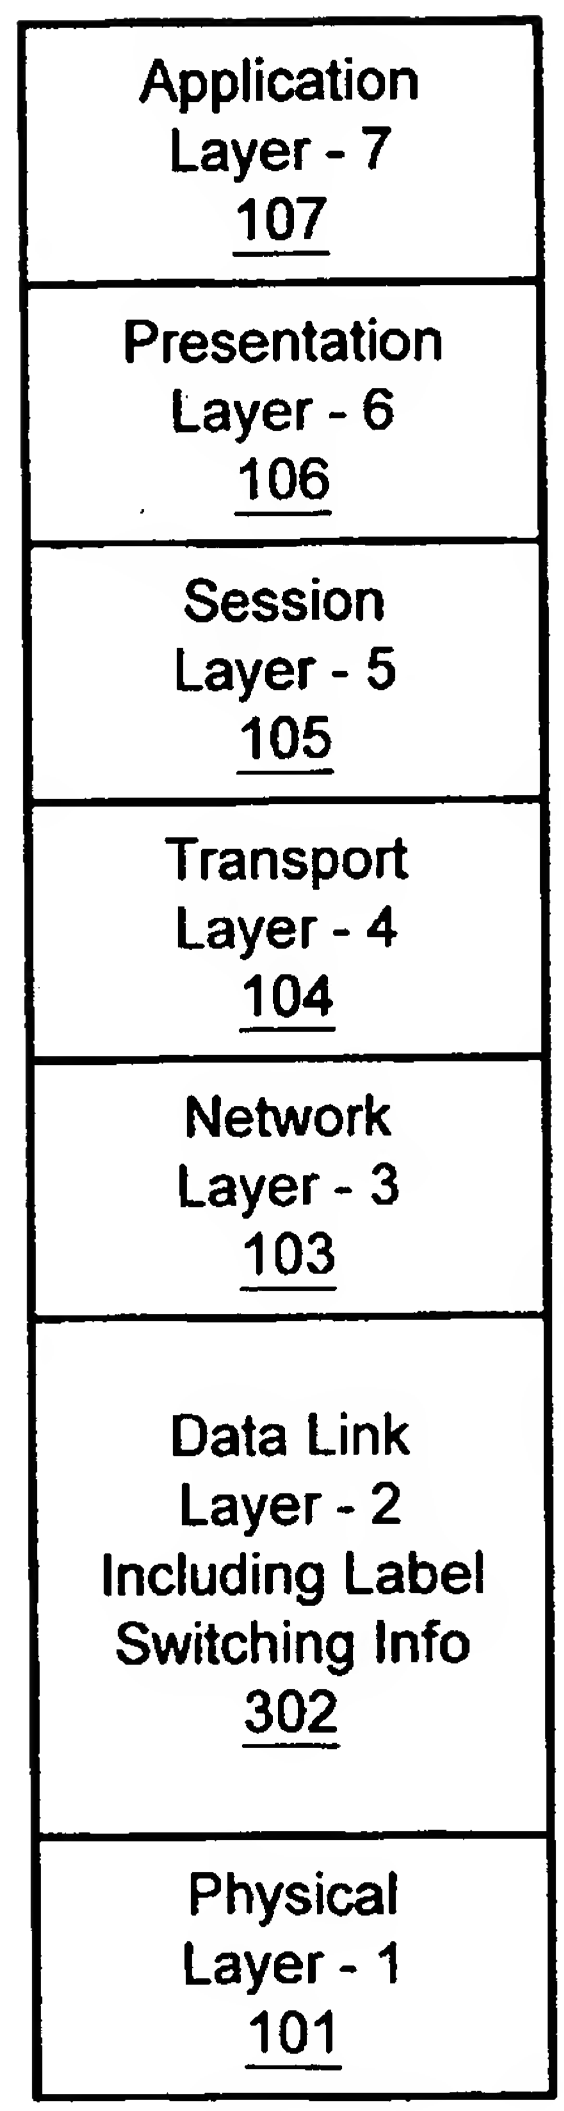

FIG. 3 - Prior Art

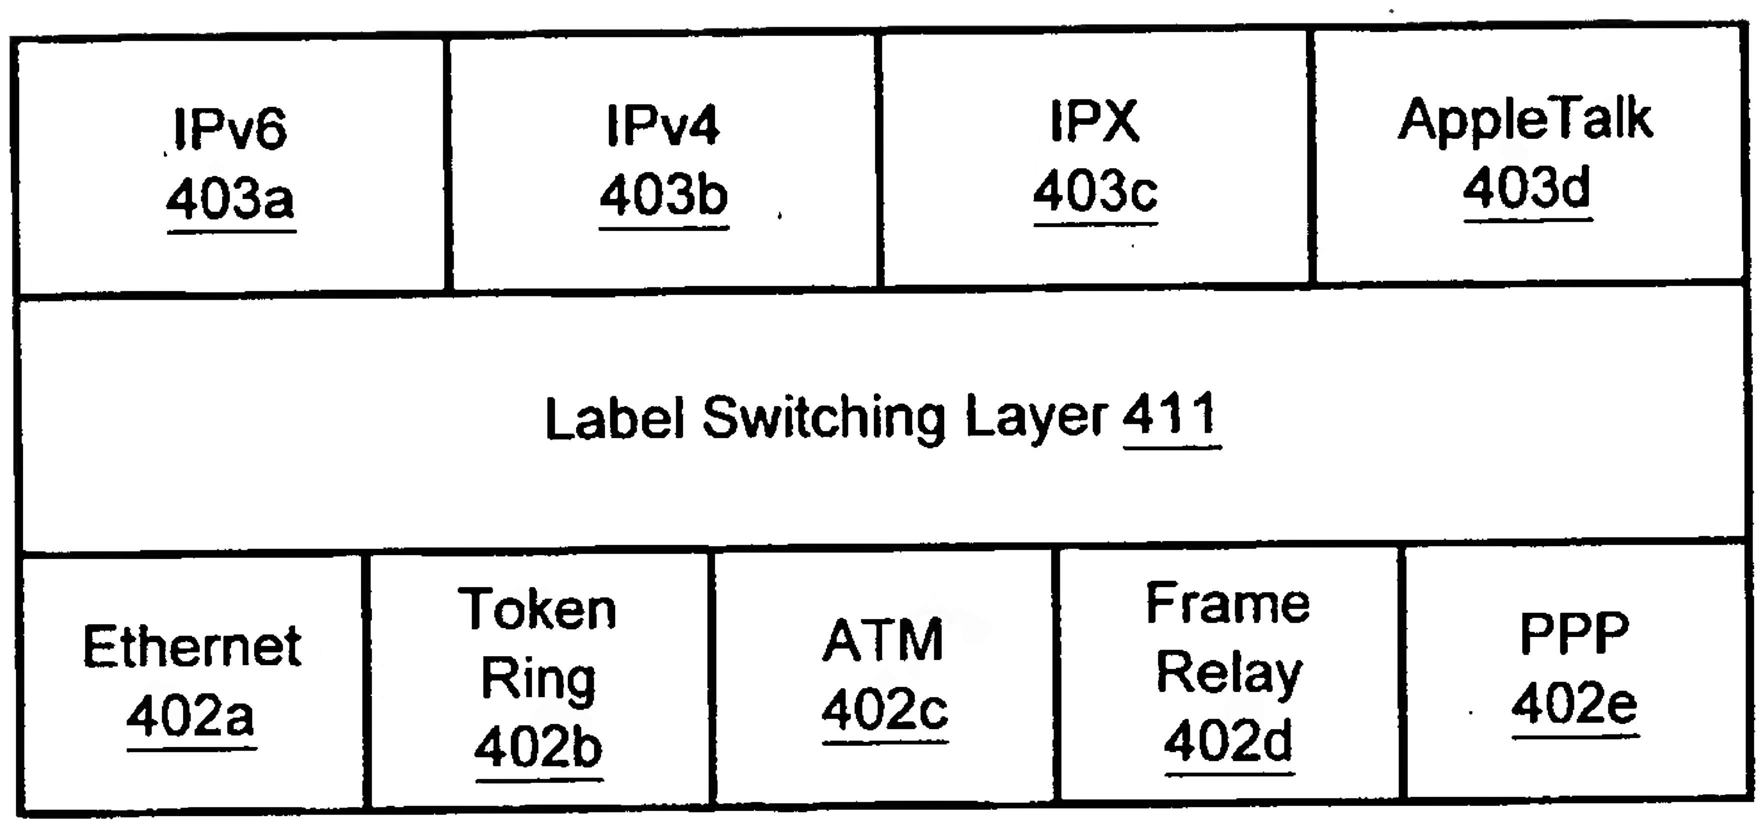

FIG. 4 - Prior Art

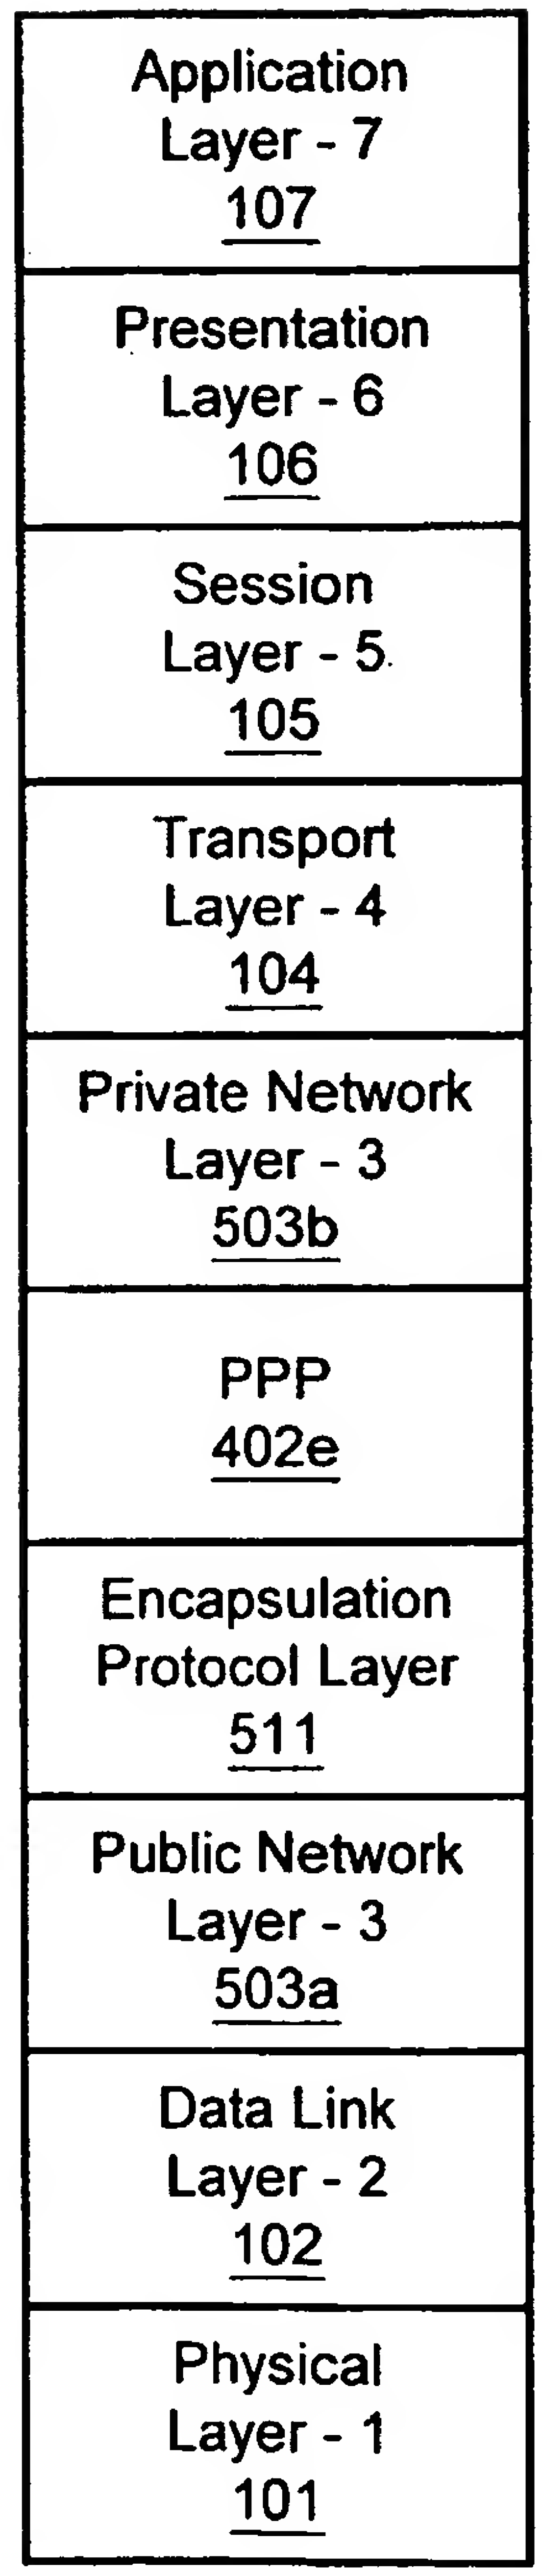

FIG. 5 - Prior Art

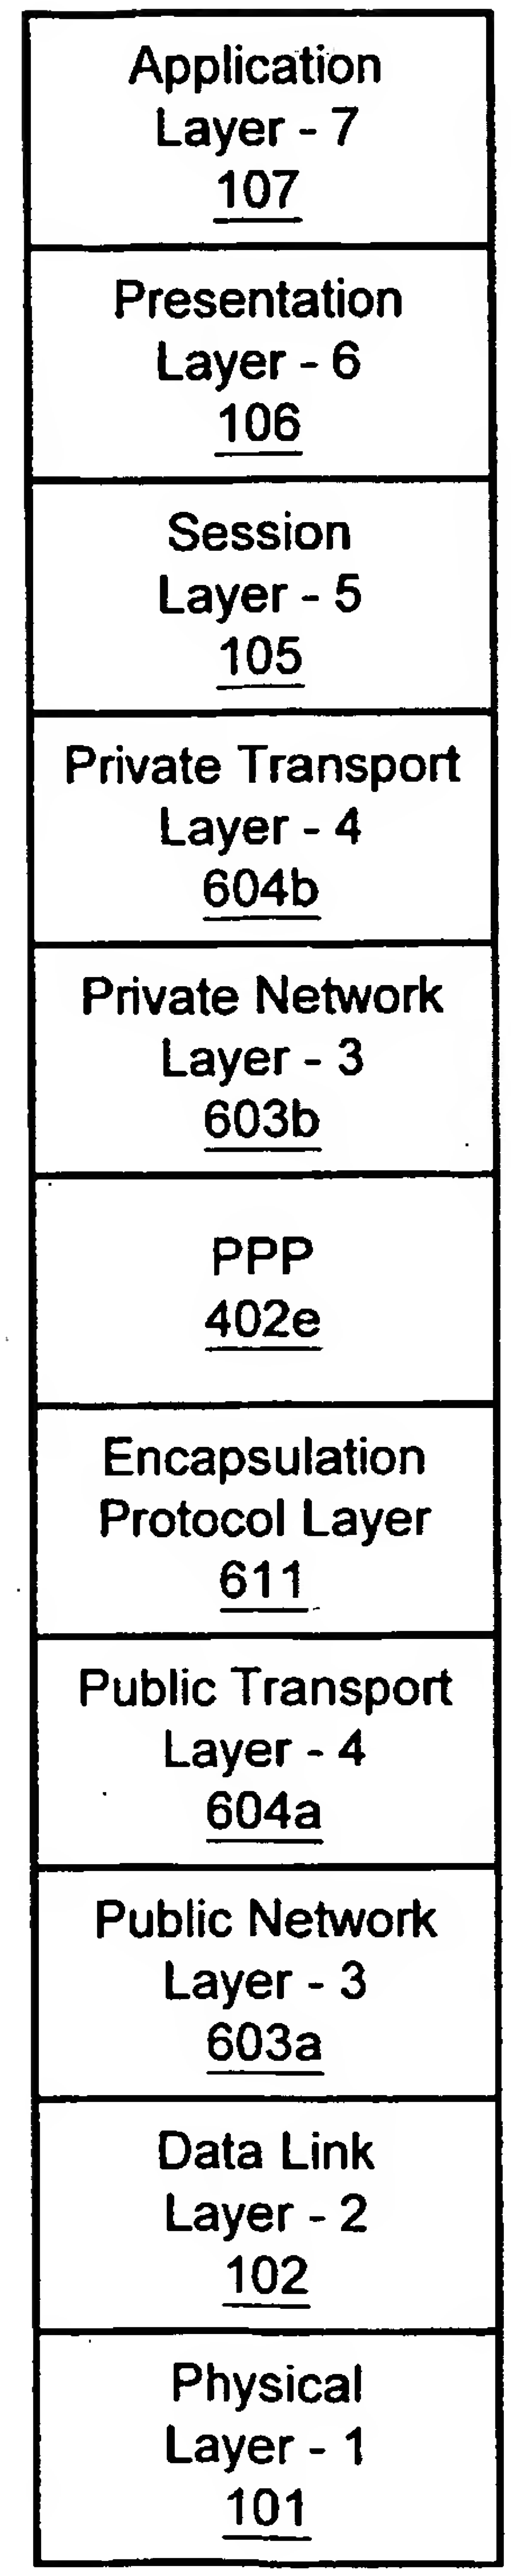

FIG. 6 - Prior Art

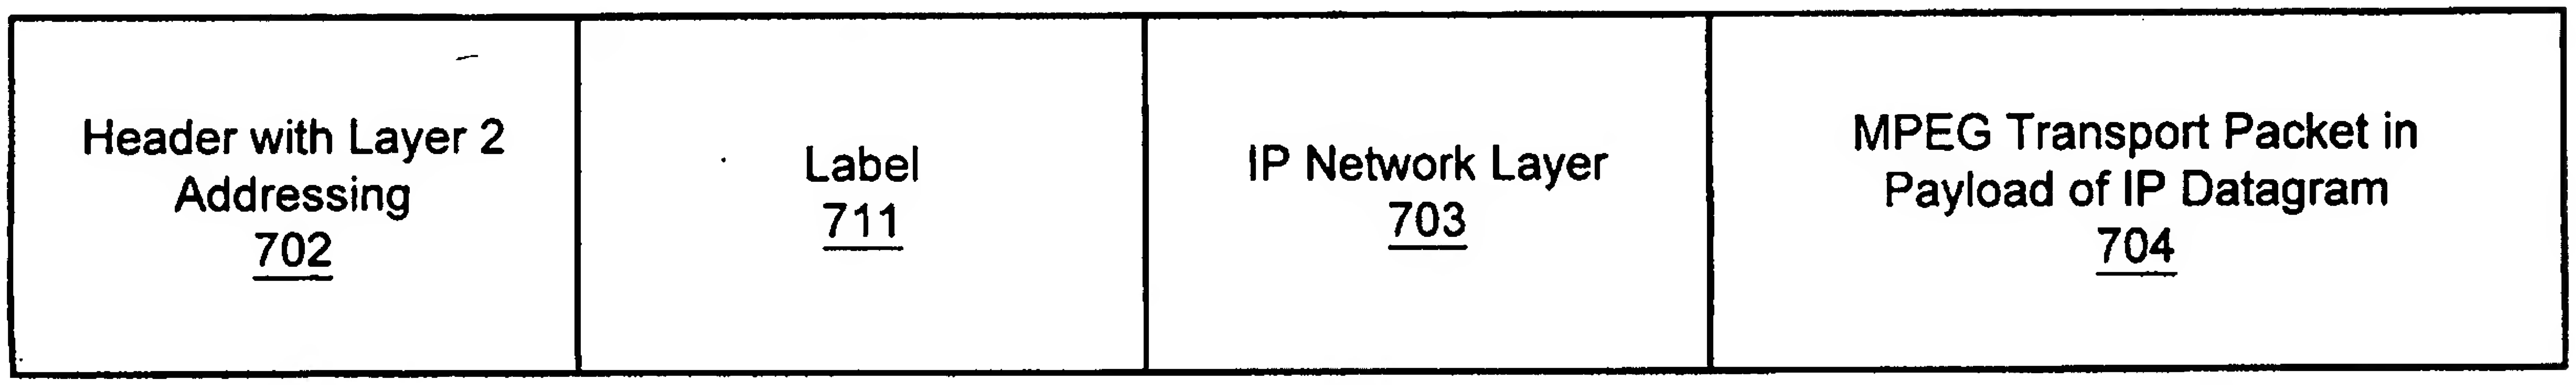

FIG. 7 - Prior Art

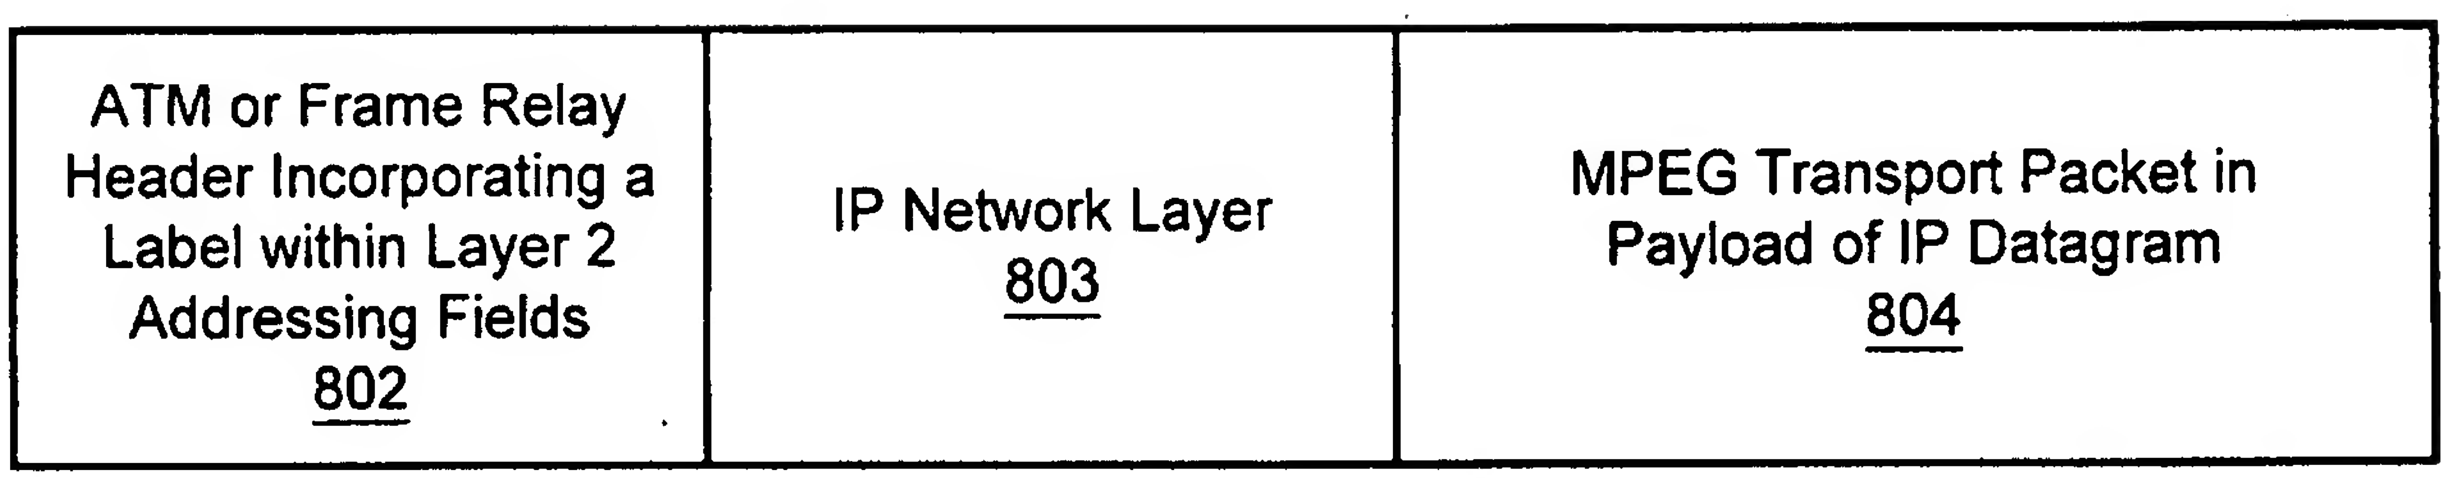

FIG. 8 - Prior Art

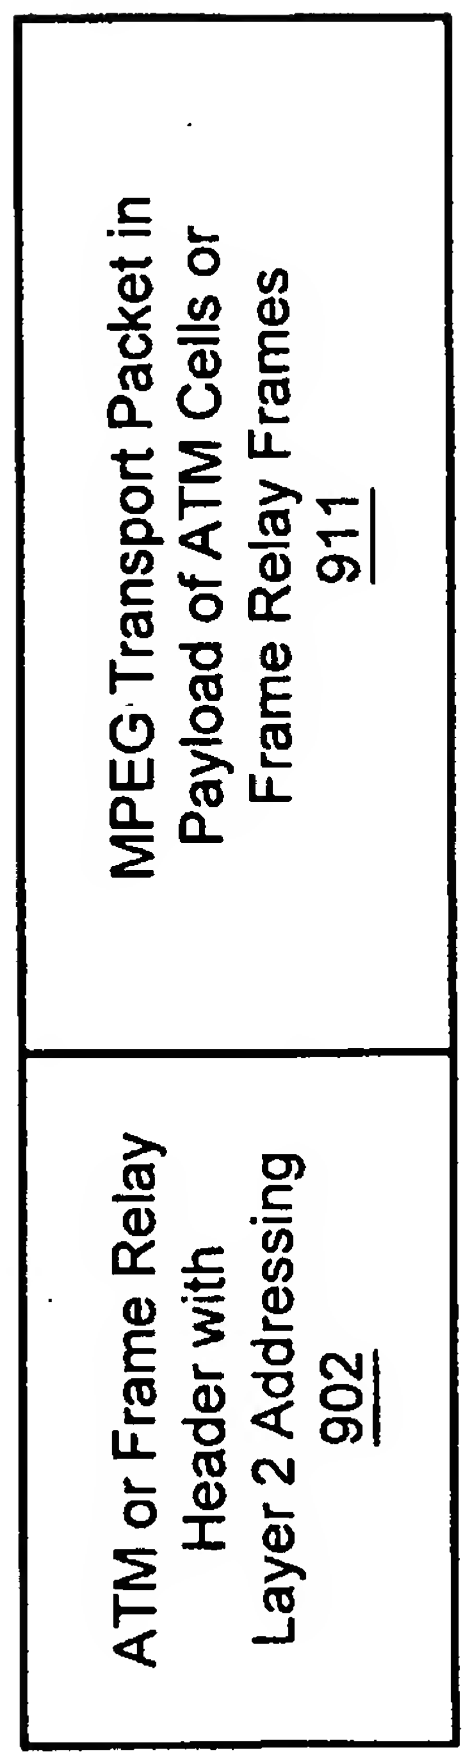

FIG. 9 - Prior Art

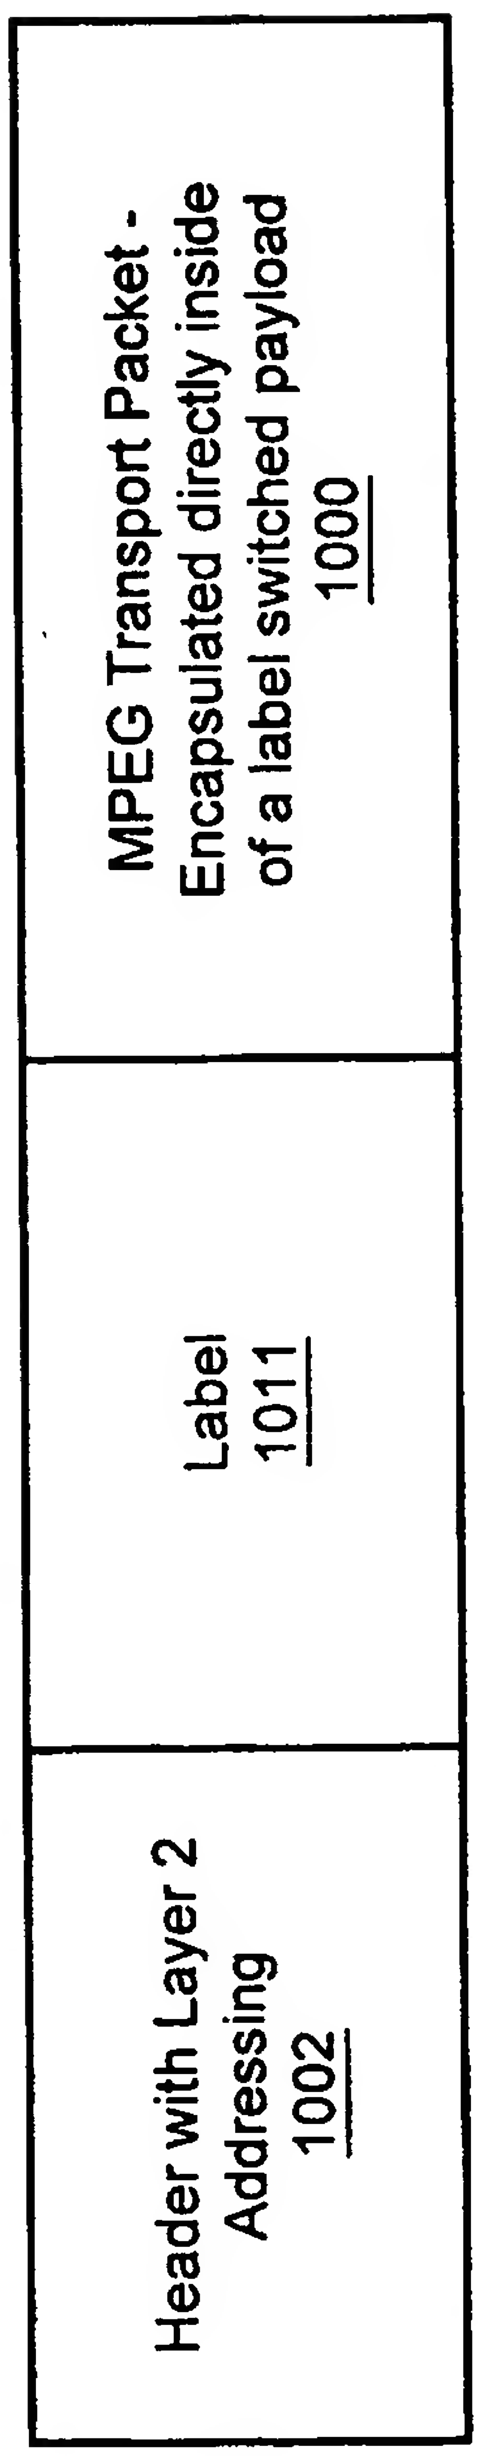
Header with Layer 2 Addressing 1002
Label 1011
MPEG Transport Packet - Encapsulated directly inside of a label switched payload 1000

FIG. 10

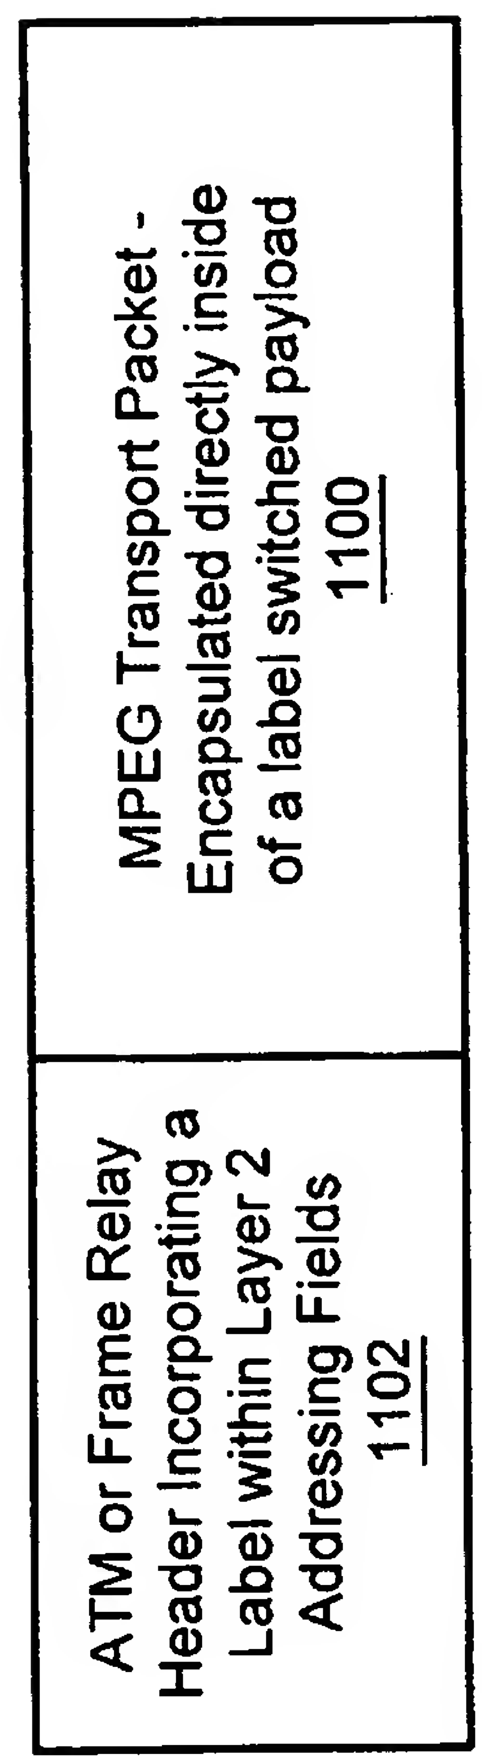
ATM or Frame Relay Header Incorporating a Label within Layer 2 Addressing Fields 1102
MPEG Transport Packet - Encapsulated directly inside of a label switched payload 1100

FIG. 11

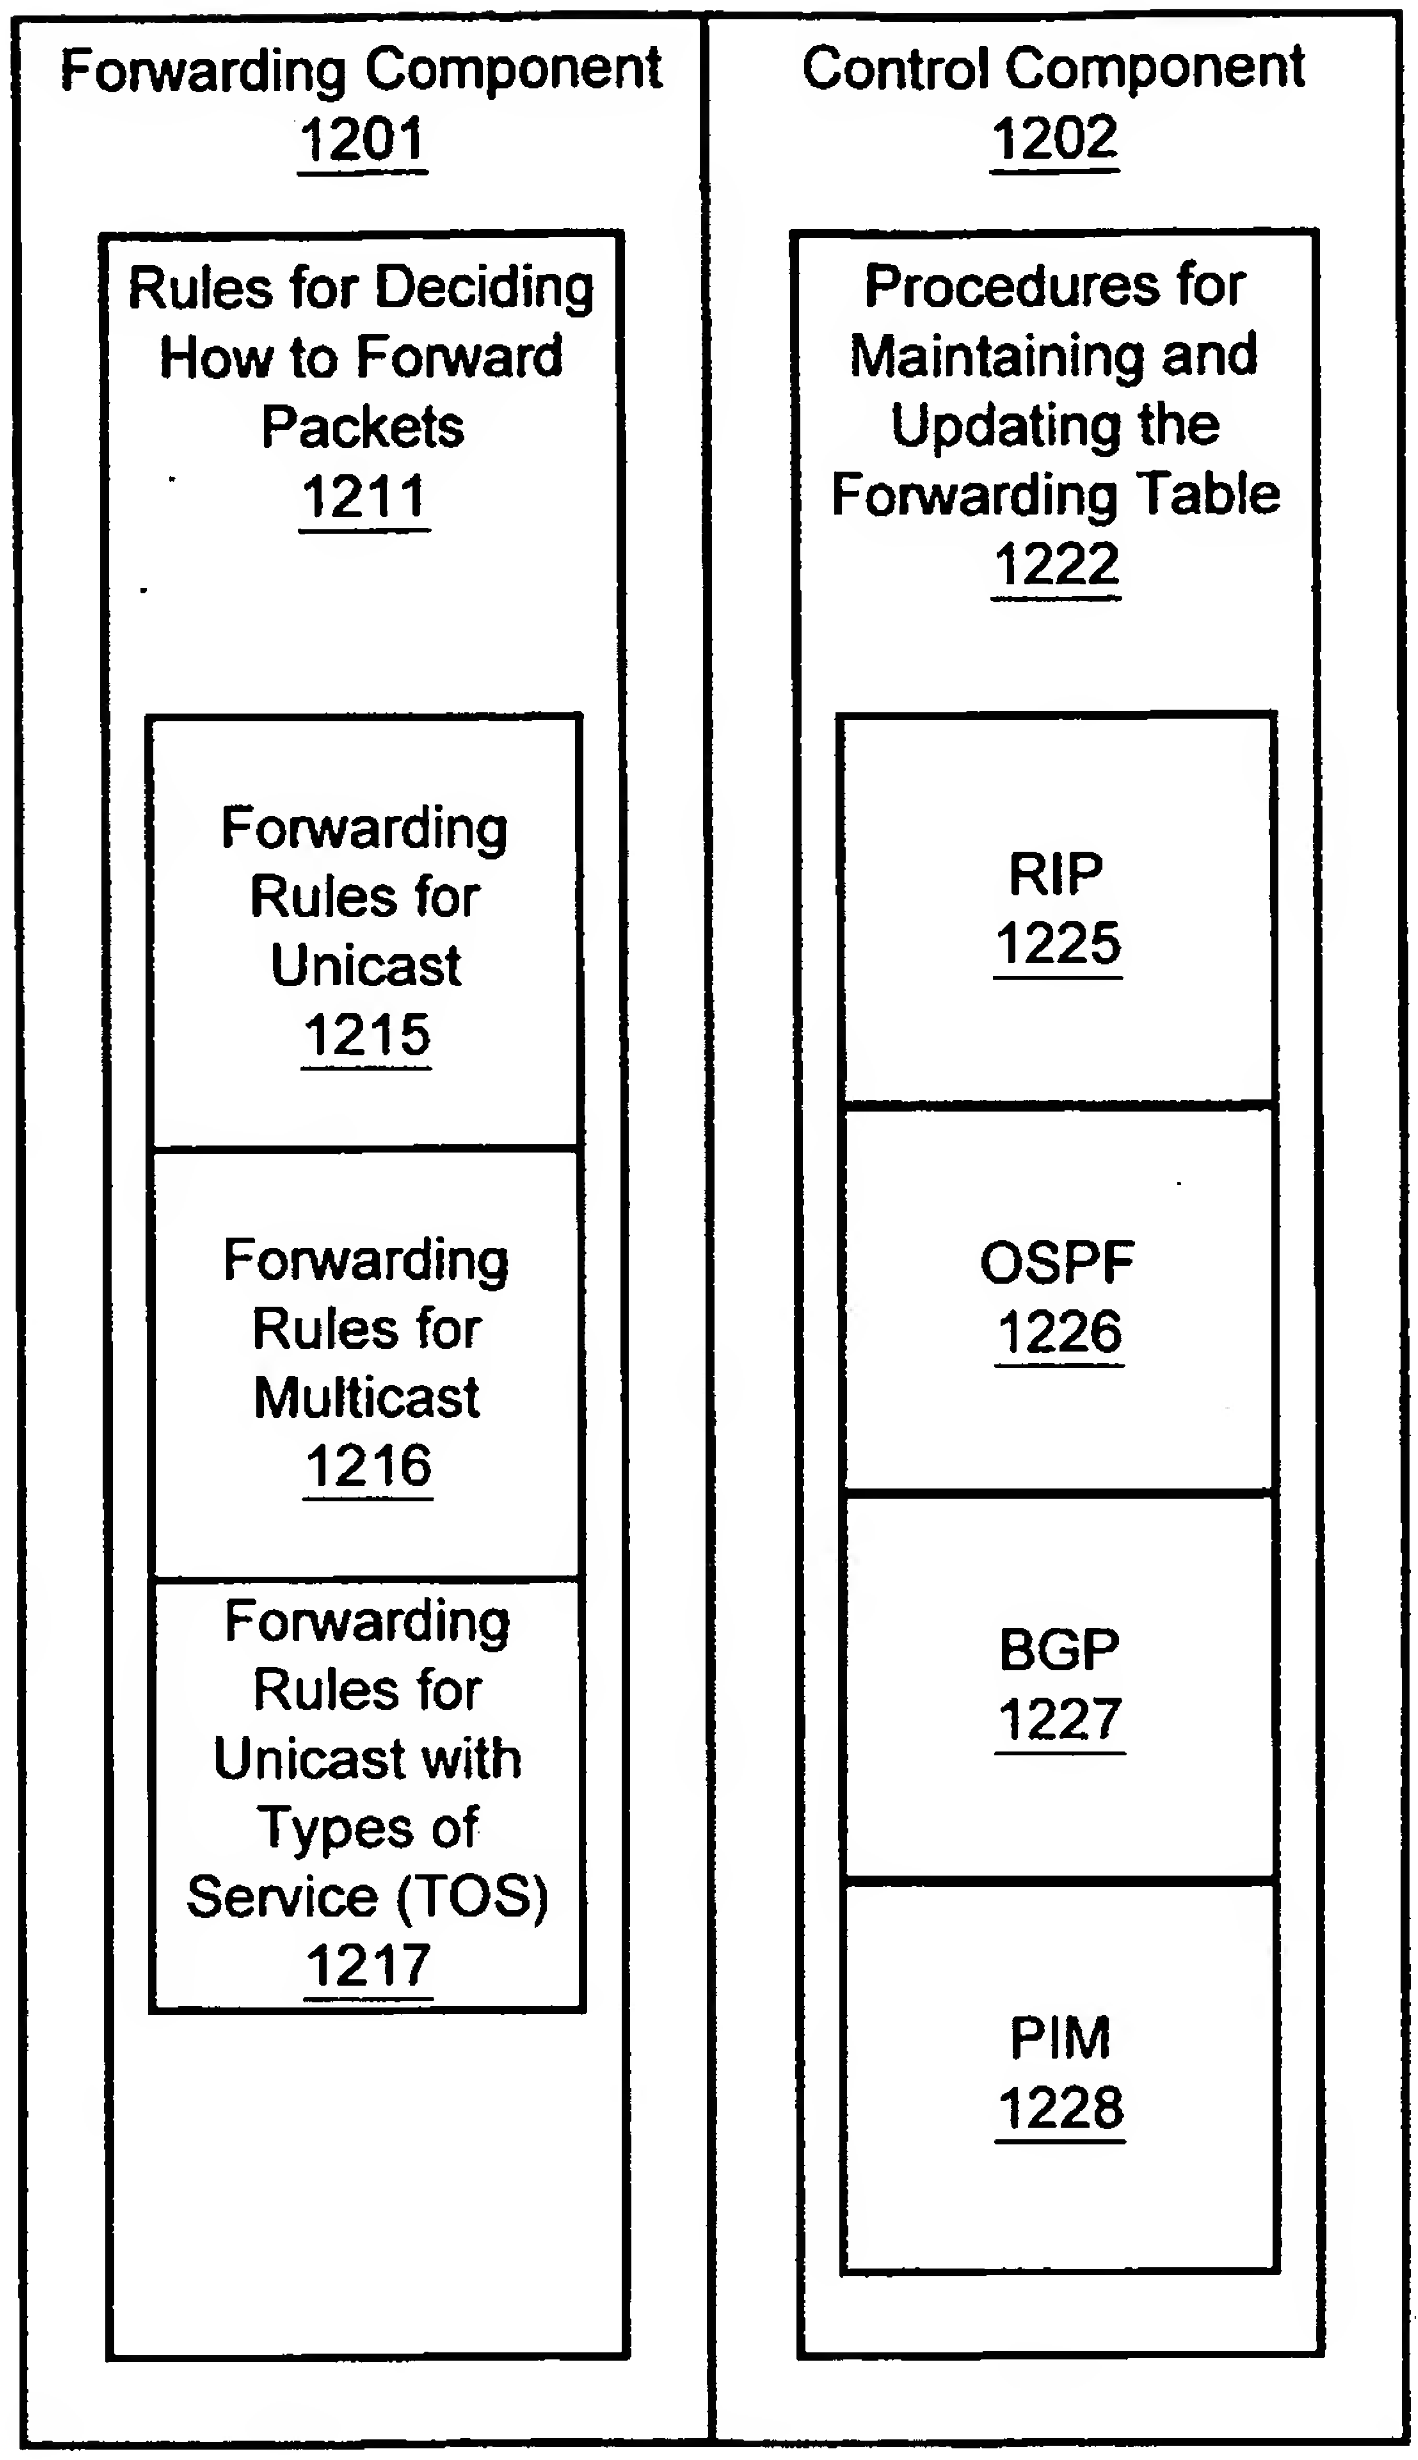

FIG. 12 - Prior Art

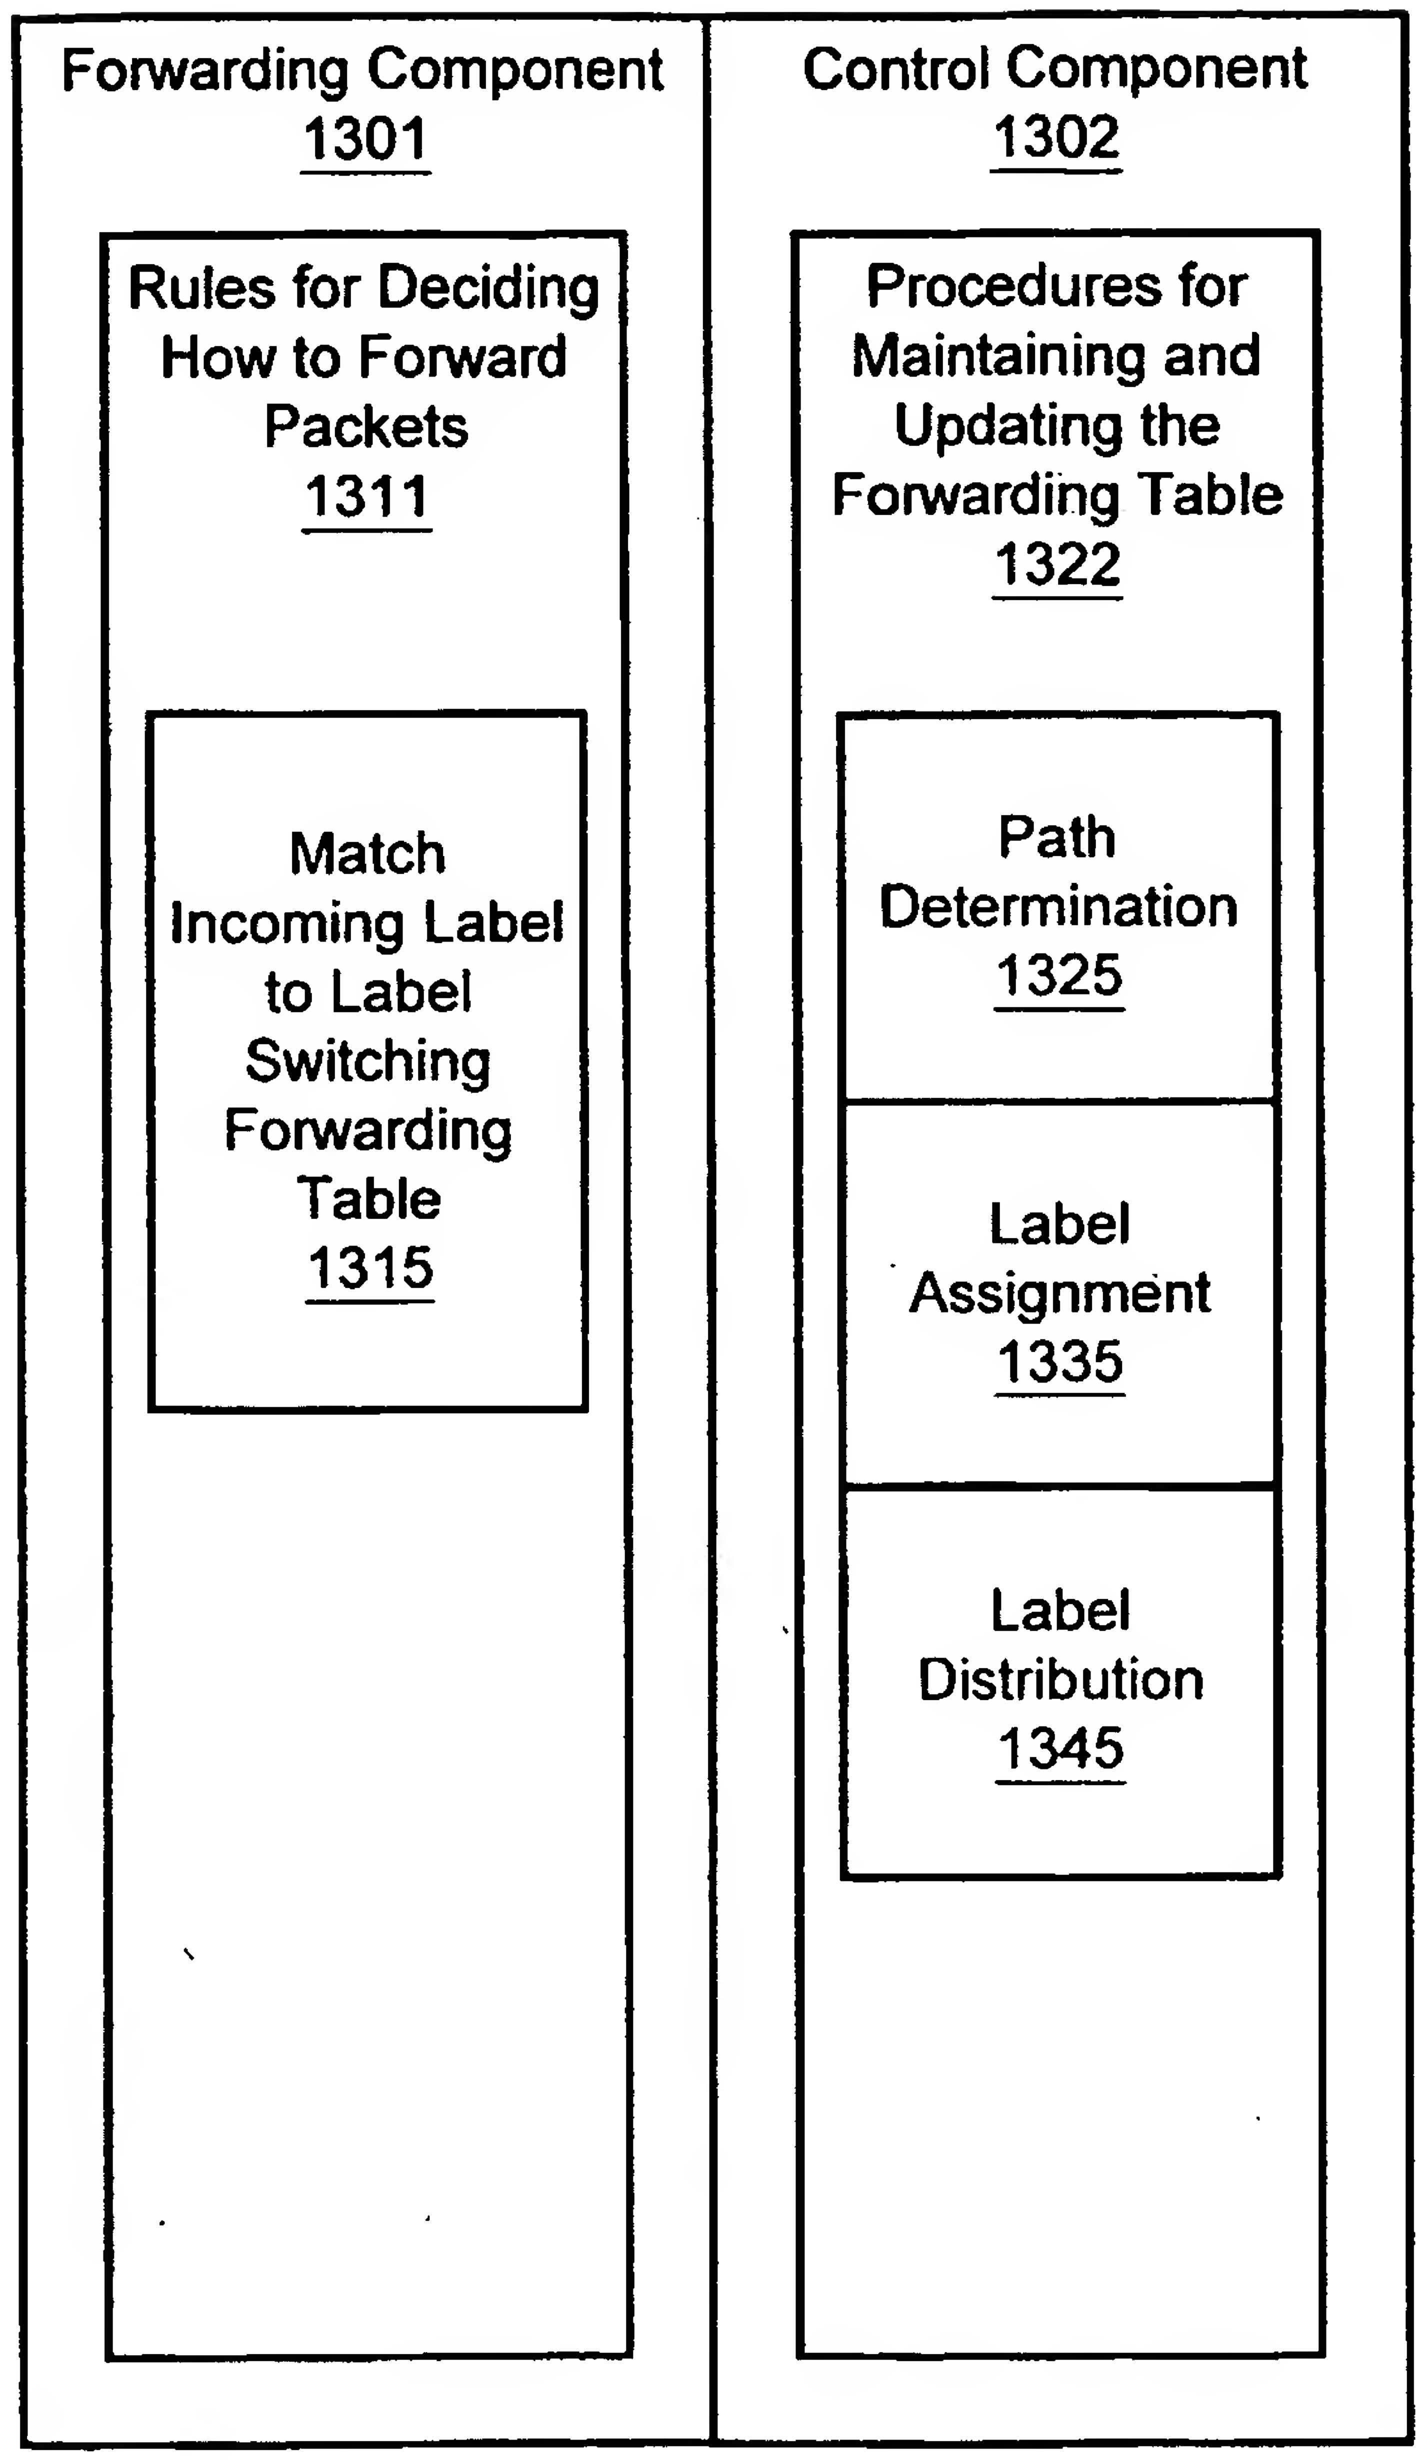
Forwarding Component
1301
Rules for Deciding How to Forward Packets
1311
Match Incoming Label to Label Switching Forwarding Table
1315
Control Component
1302
Procedures for Maintaining and Updating the Forwarding Table
1322
Path Determination
1325
Label Assignment
1335
Label Distribution
1345

FIG. 13

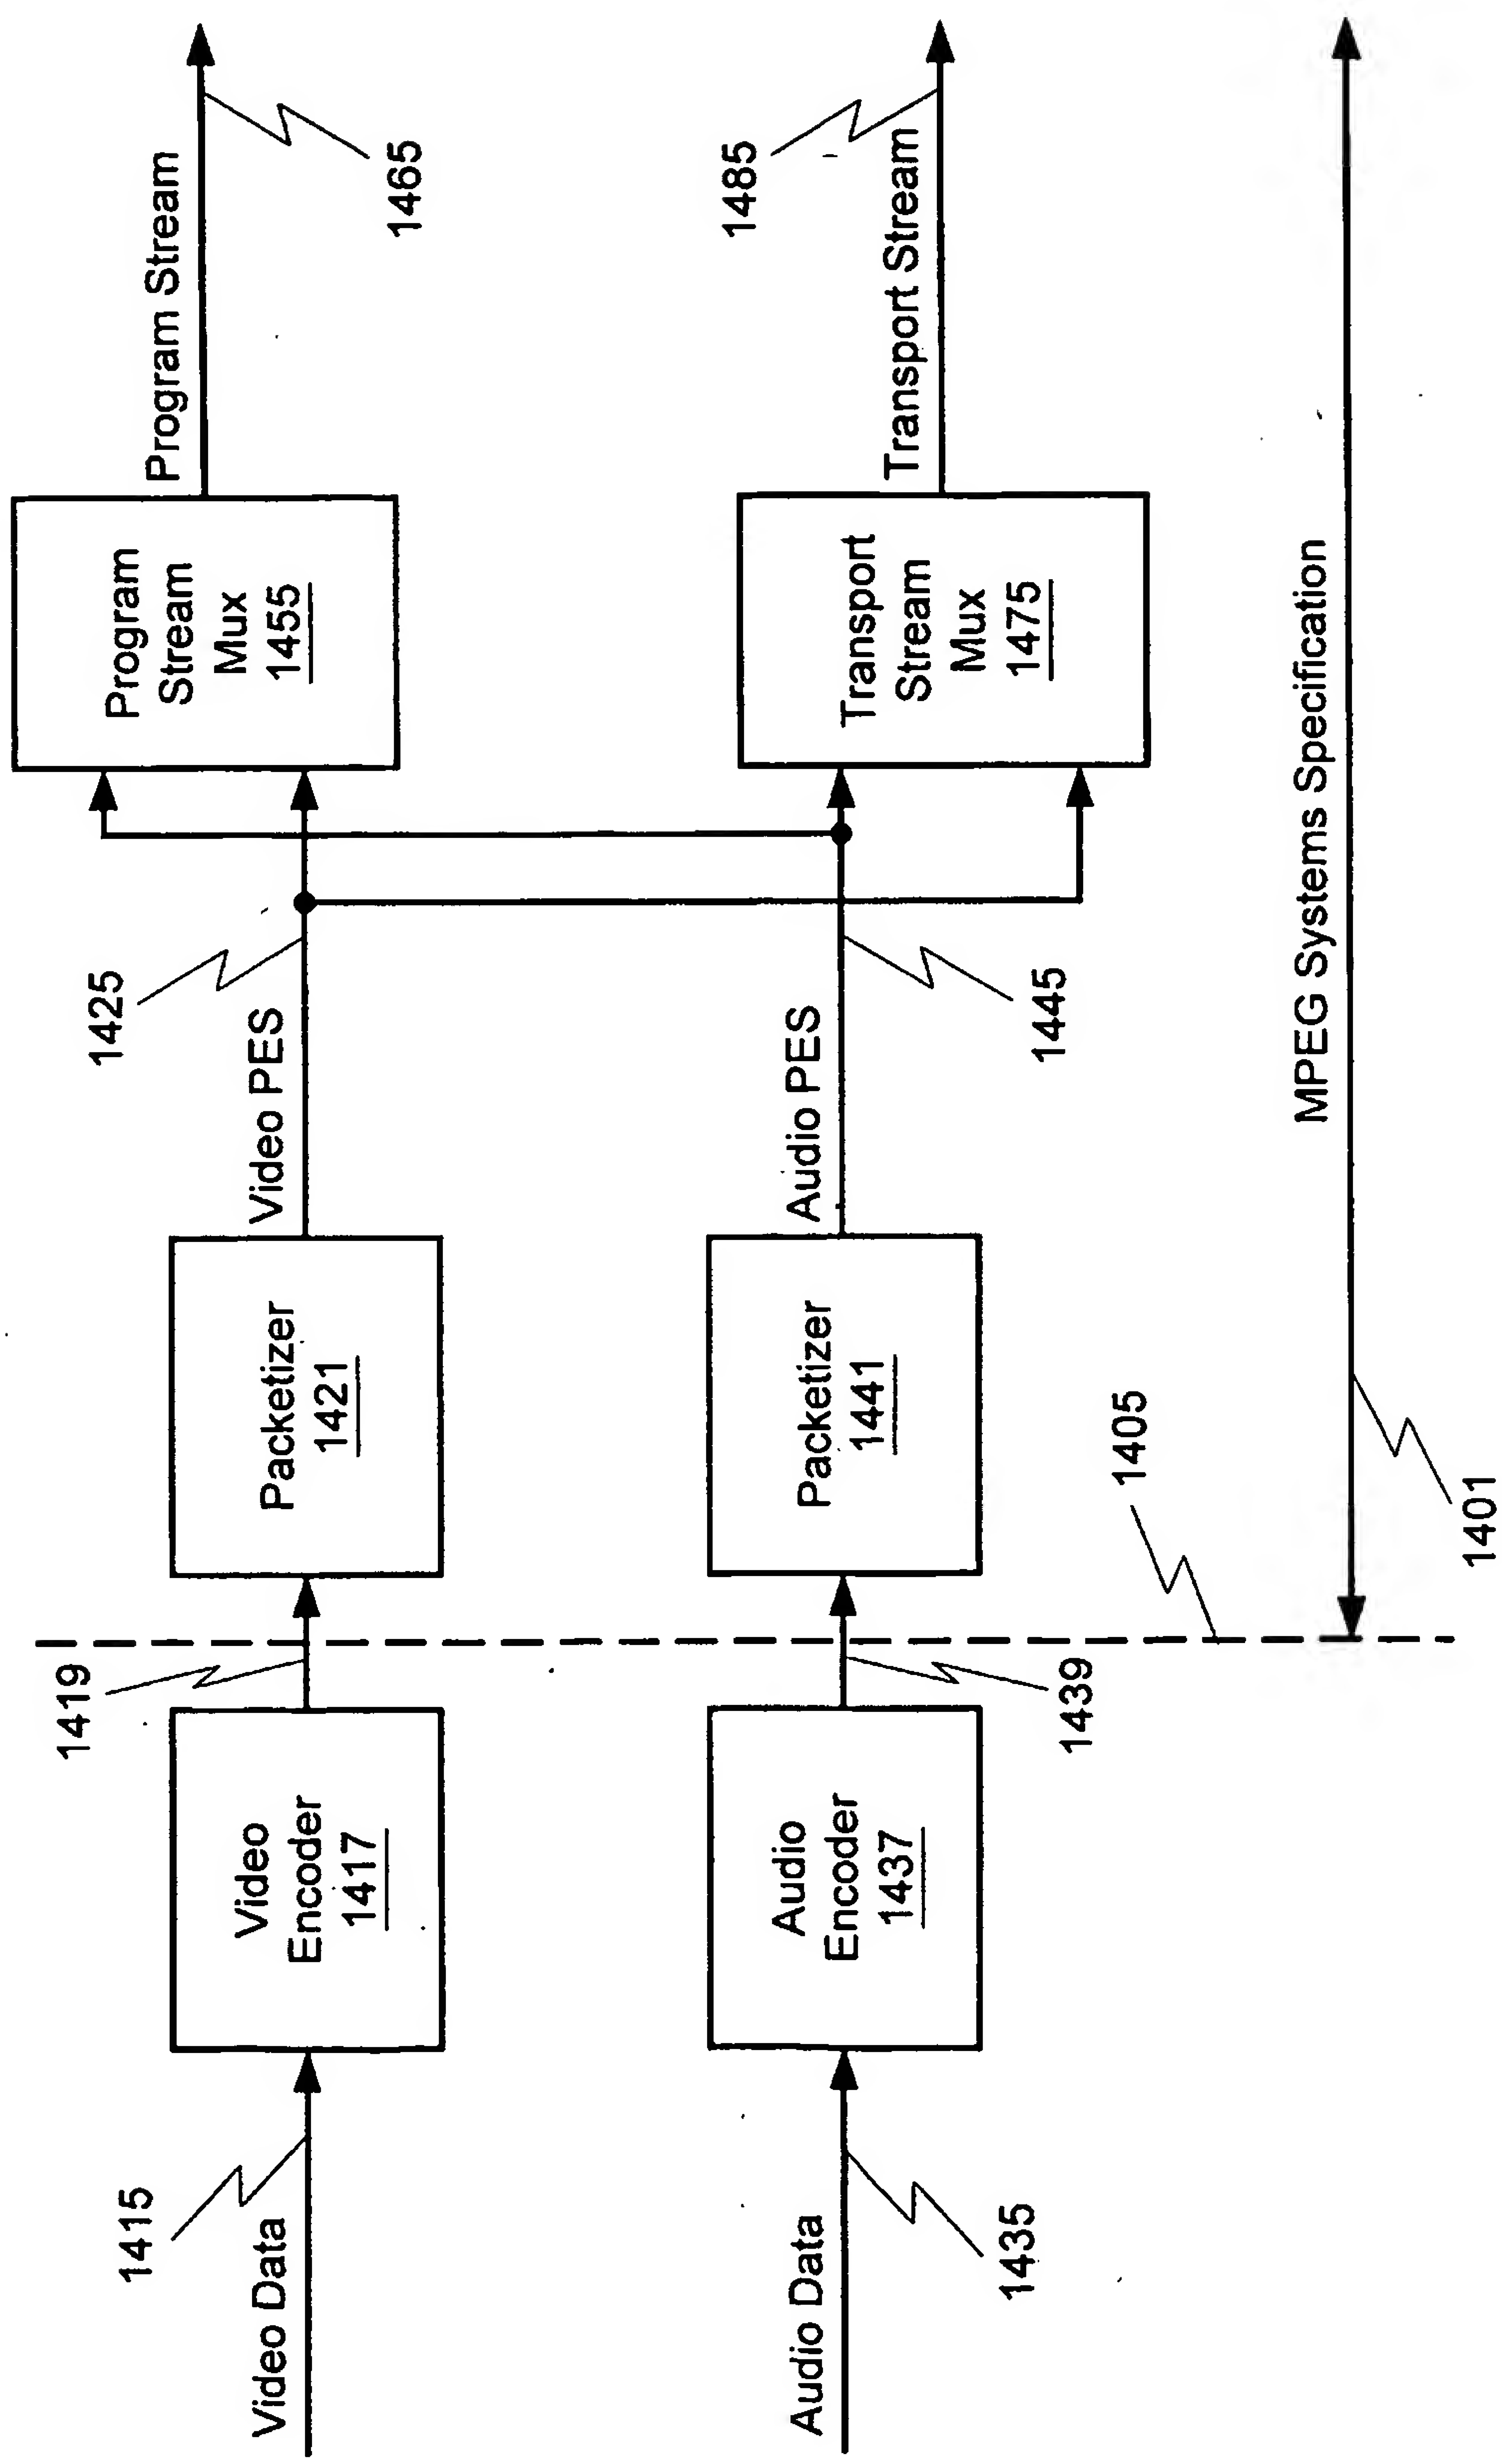

FIG. 14 - Prior Art

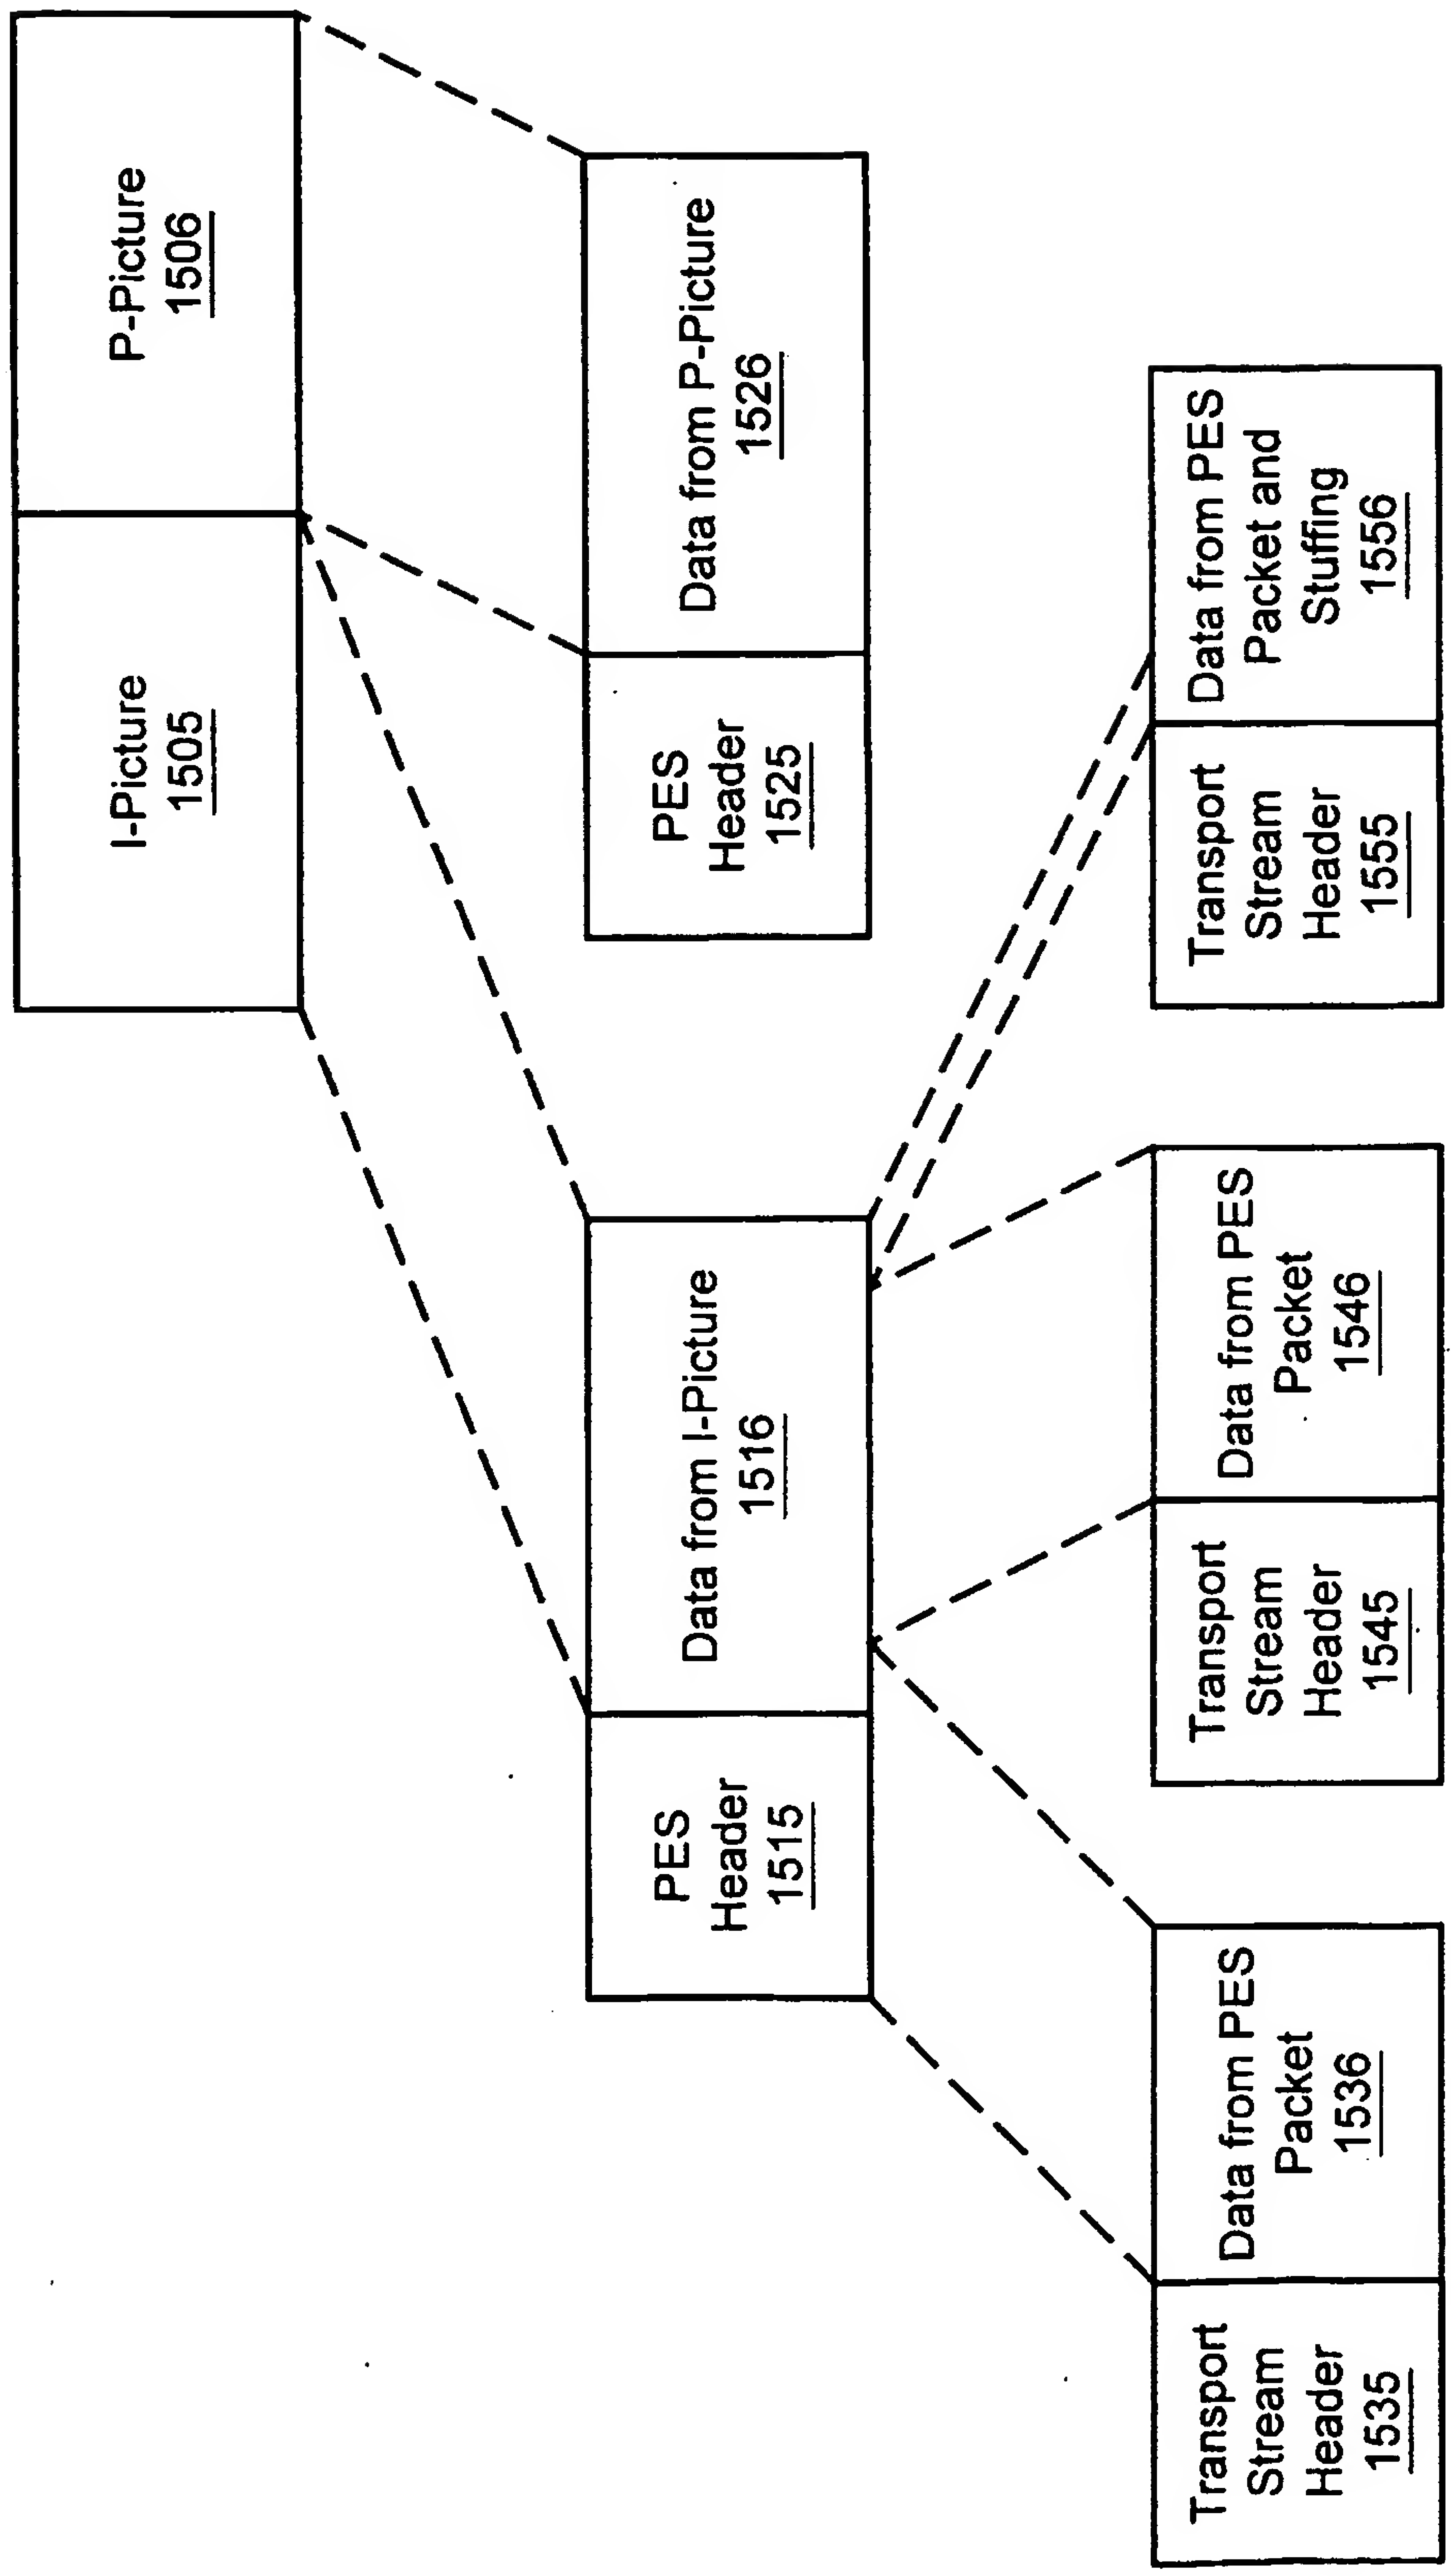

FIG. 15 - Prior Art

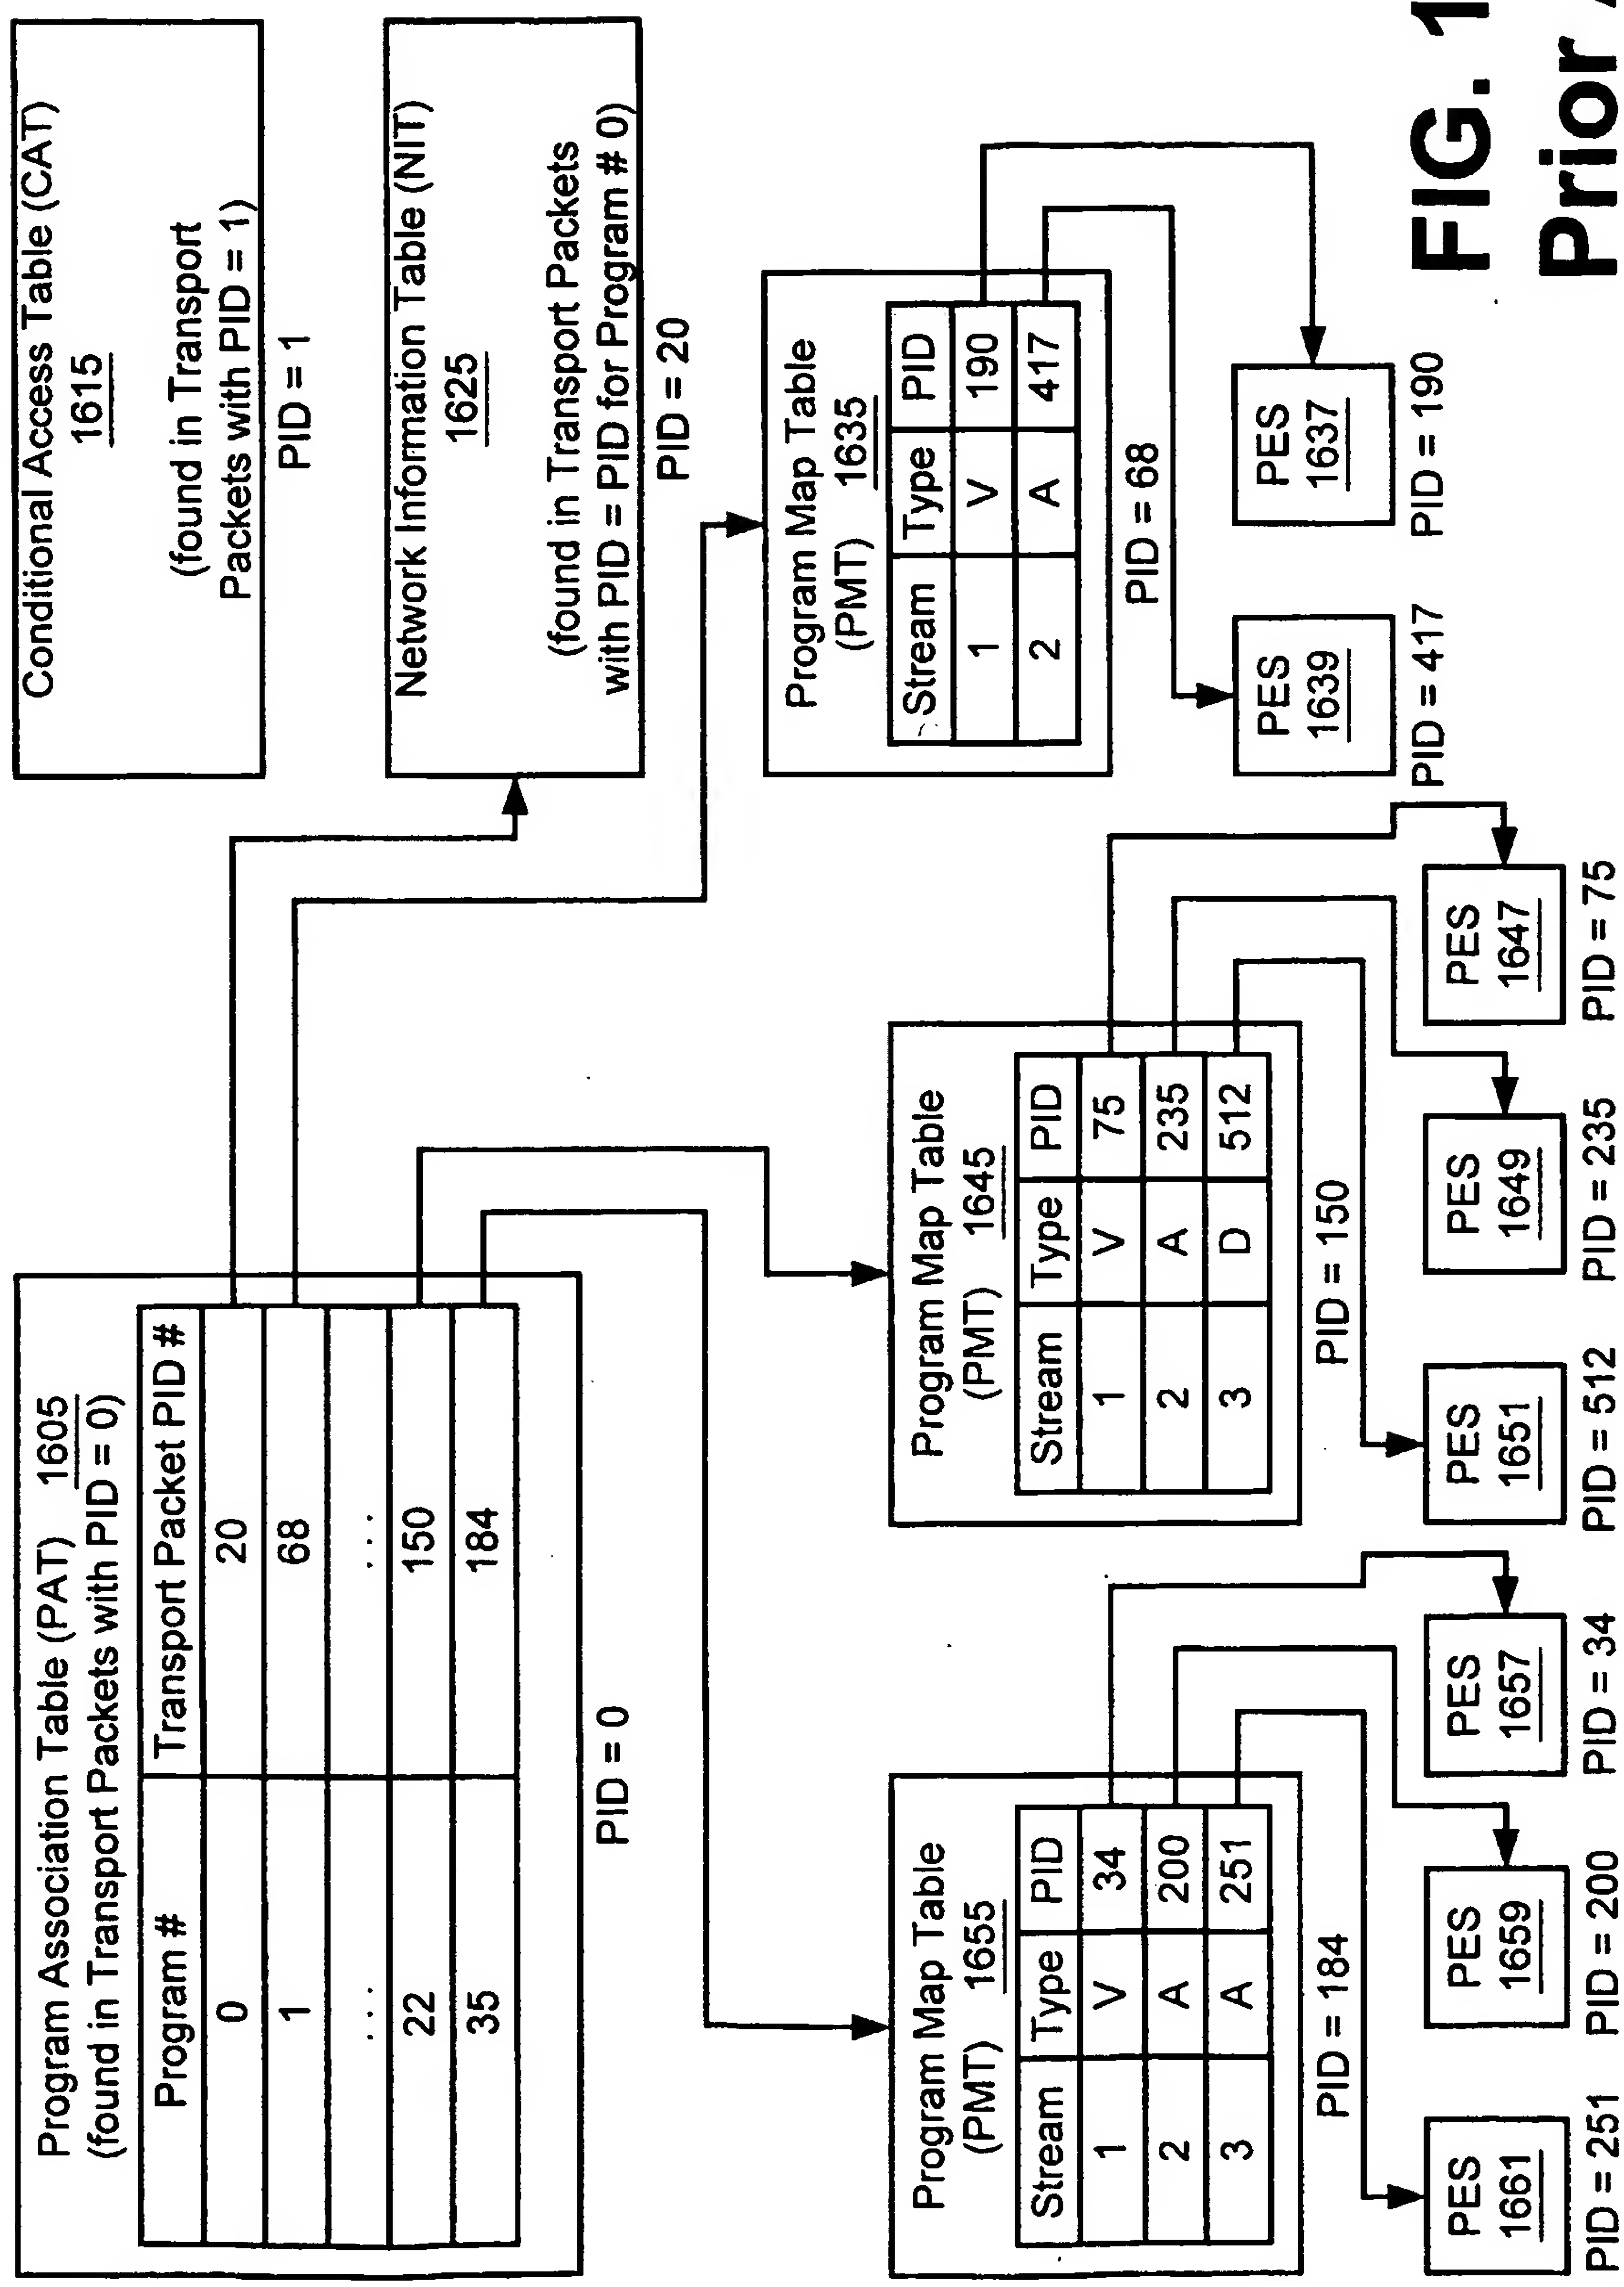

FIG. 16 - Prior Art

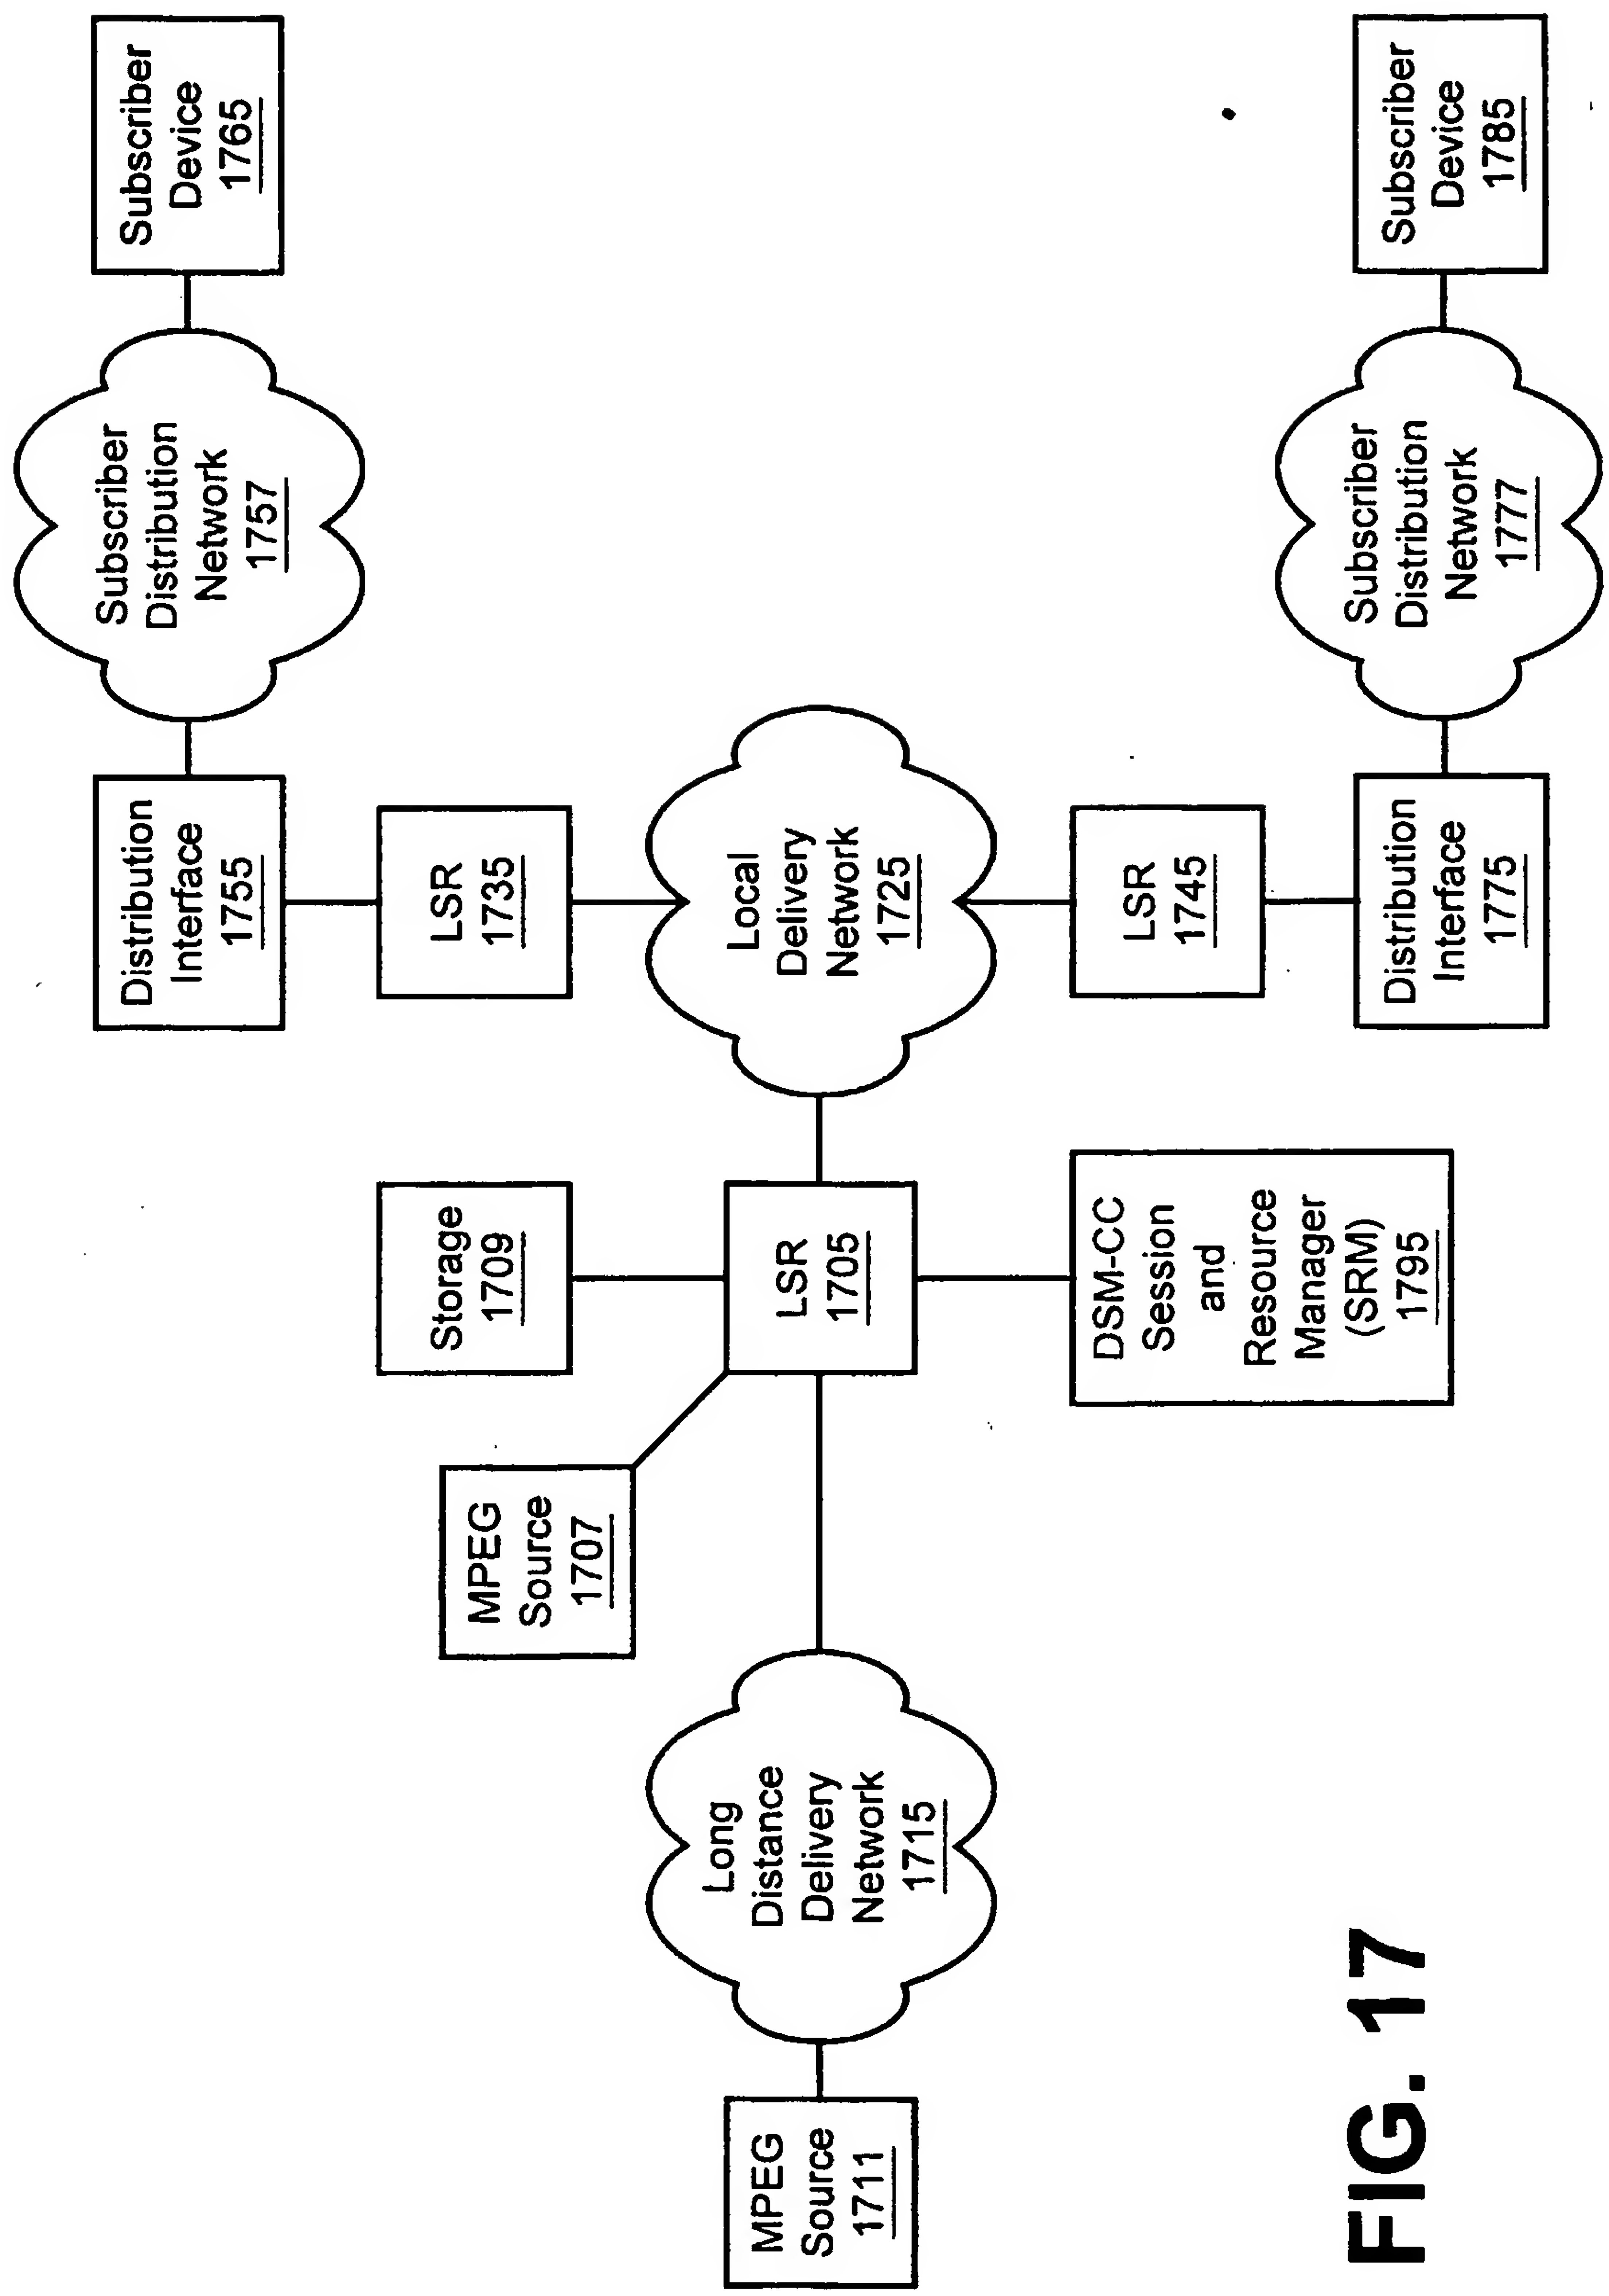
Subscriber Device 1765
Subscriber Distribution Network 1757
Distribution Interface 1755
LSR 1735
Local Delivery Network 1725
LSR 1745
Distribution Interface 1775
Subscriber Distribution Network 1777
Subscriber Device 1785
Storage 1709
LSR 1705
DSM-CC Session and Resource Manager (SRM) 1795
MPEG Source 1707
Long Distance Delivery Network 1715
MPEG Source 1711

FIG. 17